Incremental Query Processing on Big Data Streams

Leonidas Fegaras
University of Texas at Arlington
fegaras@cse.uta.edu

arXiv:1511.07846v3 [cs.DB] 6 Mar 2016

ABSTRACT

This paper addresses online query processing for large-scale, incremental data analysis on a distributed stream processing engine (DSPE). Our goal is to convert any SQL-like query to an incremental DSPE program automatically. In contrast to other approaches, we derive incremental programs that return accurate results, not approximate answers. This is accomplished by retaining a minimal state during the query evaluation lifetime and by using incremental evaluation techniques to return an accurate snapshot answer at each time interval that depends on the current state and the latest batches of data. Our methods can handle many forms of queries on nested data collections, including iterative and nested queries, group-by with aggregation, and equi-joins. Finally, we report on a prototype implementation of our framework, called MRQL Streaming, running on top of Spark and we experimentally validate the effectiveness of our methods.

Keywords: Incremental Data Processing, Distributed Stream Processing, Big Data, MRQL, Spark.

1. INTRODUCTION

In recent years, large volumes of data are being generated, analyzed, and used at an unprecedented scale and rate. Data analysis tools that process these data are typically batch programs that need to work on the complete datasets, thus repeating the computation on existing data when new data are added to the datasets. Consequently, batch processing may be prohibitively expensive for Big Data that change frequently. For example, the Web is evolving at an enormous rate with new Web pages, content, and links added daily. Web graph analysis tools, such as PageRank, which are used extensively by search engines, need to recompute their Web graph measures very frequently since they become outdated very fast. There is a recent interest in incremental Big Data analysis, where data are analyzed in incremental fashion, so that existing results on current data are reused and merged with the results of processing the new data. Incremental data processing can generally achieve better performance and may require less memory than batch processing for many data analysis tasks. It can also be used for analyzing Big Data incrementally, in batches that can fit in memory. Consequently, incremental data processing can also be useful to stream-based applications that need to process continuous streams of data in real-time with low latency, which is not feasible with existing batch analysis tools. For example, the Map-Reduce framework [14], which was designed for batch processing, is ill-suited for certain Big Data workloads, such as real-time analytics, continuous queries, and iterative algorithms. New alternative frameworks have emerged that address the inherent limitations of the Map-Reduce model and perform better for a wider spectrum of workloads. Currently, among them, the most promising frameworks that seem to be good alternatives to Map-Reduce while addressing its drawbacks are Google's Pregel [29], Apache Spark [40], and Apache Flink [20], which are in-memory distributed computing systems. There are also quite a few emerging distributed stream processing engines (DSPEs) that realize online, low-latency data processing with a series of batch computations at small time intervals, using a continuous streaming system that processes data as they arrive and emits continuous results. To cope with blocking operations and unbounded memory requirements, some of these systems build on the well-established research on data streaming based on sliding windows and incremental operators [4], which includes systems such as Aurora [1] and Telegraph [11], often yielding approximate answers, rather than accurate results. Currently, among these DSPEs, the most popular platforms are Twitter's (now Apache) Storm [35], Spark's D-Streams [46], Flink Streaming [20], Apache S4 [42], and Apache Samza [39].

This paper addresses online processing for large-scale, incremental computations on a distributed processing platform. Our goal is to convert any batch data analysis program to an incremental distributed stream processing (DSP) program automatically, without requiring the user to modify this program. We are interested in deriving DSP programs that produce accurate results, rather than approximate answers. To accomplish this task, we would need to carefully analyze the program to identify those parts that can be used to process the incremental batches of data and those parts that can be used to merge the current results with the new results of processing the incremental batches. Such analysis is hard to attain for programs written in an algorithmic programming language but can become more tractable if it is performed on declarative queries. Fortunately, most

programmers already prefer to use a higher-level query language, such as Apache Hive [26], to code their distributed data analysis applications, instead of coding them directly in an algorithmic language, such as Java. For instance, Hive is used for over 90% of Facebook Map-Reduce jobs. There are many reasons why programmers prefer query languages. First, it is hard to develop, optimize, and maintain non-trivial applications coded in a general-purpose programming language. Second, given the multitude of the new emerging distributed processing frameworks, such as Spark and Flink, it is hard to tell which one of them will prevail in the near future. Data intensive applications that have been coded in one of these paradigms may have to be rewritten as technologies evolve. Hence, it would be highly desirable to express these applications in a declarative query language that is independent of the underlying distributed platform. Furthermore, the execution of such queries can benefit from cost-based query optimization and automatic parallelism, thus relieving the application developers from the intricacies of Big Data analytics and distributed computing. Therefore, our goal is to convert batch SQL-like queries to incremental DSP programs. We have developed our framework on Apache MRQL [33], because it is both platform-independent and powerful enough to express complex data analysis tasks, such as PageRank, data clustering, and matrix factorization, using SQL-like syntax exclusively. Since we are interested in deriving incremental programs that return accurate results, not approximate answers, our focus is in retaining a minimal state during the query evaluation lifetime and across iterations, and use incremental evaluation techniques to return an accurate snapshot answer at each time interval that depends on the current state and the latest batches of data.

We have developed general, sound methods to transform batch queries to incremental queries. The first step in our approach is to transform a query so that it propagates the join and group-by keys to the query output. This is known as lineage tracking ([5, 6, 13]). That way, the values in the query output are grouped by a key combination, which corresponds the join/group-by keys used in deriving these values during query evaluation. If we also group the new data in the same way, then computations on current data can be combined with the computations on the new data by joining the data on these keys. This approach requires that we can combine computations on data that have the same lineage to derive incremental results. In our framework, this is accomplished by transforming a query to a 'monoid homomorphism' by extracting the non-homomorphic parts of the query outwards, using algebraic transformation rules, and combine them to form an answer function, which is detached from the rest of the query. We have implemented our incremental processing framework using Apache MRQL [33] on top of Apache Spark Streaming [46], which is an in-memory distributed stream processing platform. Our system is called *Incremental MRQL*. MRQL is currently the best choice for implementing our framework because other query languages for data-intensive, distributed computations provide limited syntax for operating on data collections, in the form of simple relational joins and group-bys, and cannot express complex data analysis tasks, such as PageRank, data clustering, and matrix factorization, using SQL-like syntax exclusively. Our framework though can be easily adapted to apply to other query languages, such as SQL, XQuery, Jaql, and Hive.

The contribution of this work can be summarized as follows:

- We present a general automated method to convert most distributed data analysis queries to incremental stream processing programs.
- Our methods can handle many forms of queries, including iterative and nested queries, group-by with aggregation, and joins on one-to-many relationships.
- We have extended our framework to handle deletions by using state increments to diminish the state in such a way that the query results would be the same as those we could have gotten on current data after deletion.
- We report on a prototype implementation of our framework using Apache MRQL running on top of Apache Spark Streaming. We show the effectiveness of our method through experiments on four queries: groupBy, join-groupBy, k-means clustering, and PageRank.

The rest of this paper is organized as follows. Section 2 illustrates our approach through examples. Section 3 compares our work with related work. Section 4 describes our earlier work on MRQL query processing. Section 5 defines the algebraic operators used in the MRQL algebra. Section 6 presents transformation rules for converting algebraic terms to a normal form that is easier to analyze. Section 7 describes an inference algorithm that statically infers the merge function that merges the current results with the results of processing the new data. Section 8 describes our method for transforming algebraic terms to propagate all keys used in joins and group-bys to the query output. Section 9 describes an algorithm that pulls the non-homomorphic parts of a query outwards, deriving a homomorphism. Section 10, gives an example of the transformation of a batch query to an incremental query. Section 11 extends our framework to handle data deletion. Section 12 gives some implementation details. Section 13 presents experiments that evaluate the performance of our incremental query processing techniques for four queries.

2. HIGHLIGHTS OF OUR APPROACH

A dataset in our framework is a bag (multiset) that consists of arbitrarily complex values, which may contain nested bags and hierarchical data, such as XML and JSON fragments. We are considering continuous queries over a number of streaming data sources, S_i, for $0 < i \leq n$. A data stream S_i in our framework consists of an initial dataset, followed by a continuous stream of incremental batches ΔS_i that arrive at regular time intervals Δt. In addition to streaming data, there may be other input data sources that remain invariant through time. A streaming query in our framework can be expressed as $q(\overline{S})$, where an $S_i \in \overline{S}$ is a streaming data source. Incremental stream processing is feasible when we can derive the query results at time $t+\Delta t$ by simply combining the query results at time t (i.e., the current results) with the results of processing the incremental batches ΔS_i only, rather than the entire streams $S_i \uplus \Delta S_i$, where $\uplus$ is the additive (bag) union. This is possible if $q(\overline{S \uplus \Delta S})$ can be expressed in terms of $q(\overline{S})$ (the current query result) and $q(\overline{\Delta S})$ (the incremental query result), that is, when $q(\overline{S})$ is a homomorphism[1] over $\overline{S}$. But some queries, such as counting

[1]We use the term *homomorphism* throughout the paper as an abbreviation of *monoid homomorphism*.

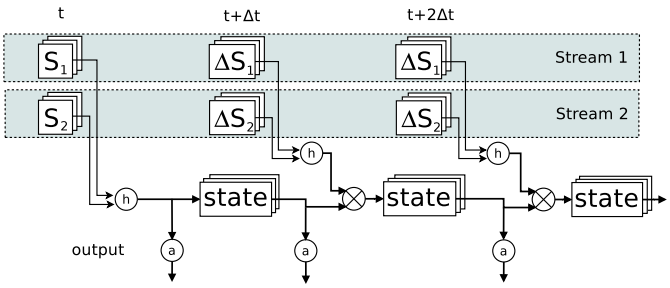

Figure 1: Incremental Query Processing

the number of distinct elements in a stream or calculating average values after a group-by, are not homomorphisms. For such queries, we break q into two functions a and h, so that $q(\overline{S}) = a(h(\overline{S}))$ and h is a homomorphism. The function h is a homomorphism if $h(\overline{S \uplus \Delta S}) = h(\overline{S}) \otimes h(\overline{\Delta S})$ for some monoid $\otimes$ (an associative function with a zero element $\otimes_z$). For example, the query that counts the number of distinct elements can be broken into the query h that returns the list of distinct elements (a homomorphism), followed by the answer query a that counts these elements. Ideally, we would like most of the computation in q to be done in h, leaving only some computationally inexpensive data mappings to the answer function a. Note that, the obvious solution when a is equal to q and h is the union of data sources, is also the worst-case scenario that we try to avoid, since it basically requires to compute the new result from the entire input, $\overline{S \uplus \Delta S}$. On the other hand, in the special case when a is the identity function and $\otimes$ is equal to $\uplus$, the output at each time interval can be simply taken to be only $h(\overline{\Delta S})$, which is the output we would expect to get from a fixed window system (i.e., new batches of output from new batches of input).

If we split q into a homomorphism h and an answer function a, then we can calculate h incrementally by storing its results into a state and then using the current state to calculate the next h result. Initially, state $= \otimes_z$ or, if there are initial stream data, state $= h(\overline{S})$. Then, at each time interval Δt, the query answer is calculated from the state, which becomes equal to $h(\overline{S \uplus \Delta S})$:

$$\text{state} \leftarrow \text{state} \otimes h(\overline{\Delta S})$$
$$\textbf{return } a(\text{state})$$

In Spark, for example, the state and the invariant data sources are stored in memory as Distributed DataSets (Spark's RDDs [45]) and are distributed across the worker nodes. However, the streaming data sources are implemented as Discretized Streams (Spark's D-Streams [46]), which are also distributed.

Our framework works best for queries whose output is considerably smaller than their input, such as for data analysis queries that aggregate data. Such queries would require a smaller state and impose less processing overhead to $\otimes$.

Figure 1 shows the evaluation of an incremental query $q(S_1, S_2) = a(h(S_1, S_2))$ over two streaming data sources, where $h(S_1 \uplus \Delta S_1, S_2 \uplus \Delta S_2) = h(S_1, S_2) \otimes h(\Delta S_1, \Delta S_2)$.

Our query processing system performs the following tasks:

1. It pulls all non-homomorphic parts of a query q out from the query, using algebraic transformations.
2. It collects these non-homomorphic parts into an answer function a, leaving an algebraic homomorphic term h, such that $q(\overline{S}) = a(h(\overline{S}))$.
3. From the homomorphic algebraic term $h(\overline{S})$, our system derives a merge function $\otimes$, such that $h(\overline{S \uplus \Delta S}) = h(\overline{S}) \otimes h(\overline{\Delta S})$.

These tasks are, in general, hard to attain for a program expressed in an algorithmic programming language. Instead, these tasks can become more tractable if they are performed on higher-order operations, such as the MRQL query algebra [18, 17]. Although all algebraic operations used in MRQL are homomorphic, their composition may not be. We have developed transformation rules to derive homomorphisms from compositions of homomorphisms, and for pulling non-homomorphic parts outside a query. Our methods can handle most forms of queries on nested data sets, including iterative queries, complex nested queries with any form and any number of nesting levels, general group-bys with aggregations, and general one-to-one and one-to-many equi-joins. Our methods cannot handle non-equi-joins and many-to-many equi-joins, as they are notoriously difficult to implement efficiently in a streaming or an incremental computing environment.

Example. For example, consider the following query (all queries are expressed in MRQL):

$$\begin{aligned} q(S_1, S_2) = \ & \textbf{select } (\mathsf{x.A}, \textbf{avg}(\mathsf{y.D})) \\ & \textbf{from } \mathsf{x} \textbf{ in } S_1, \ \mathsf{y} \textbf{ in } S_2 \\ & \textbf{where } \mathsf{x.B} = \mathsf{y.C} \\ & \textbf{group by } \mathsf{x.A} \end{aligned}$$

where S_1 and S_2 in $q(S_1, S_2)$ are streaming data sources. Unfortunately, $q(S_1, S_2)$ is not a homomorphism over S_1 and S_2, that is, $q(S_1 \uplus \Delta S_1, S_2 \uplus \Delta S_2)$ cannot be expressed in terms of $q(S_1, S_2)$ and $q(\Delta S_1, \Delta S_2)$ exclusively. The intrinsic reason behind this is that there is no lineage in the query output that links a pair in the query result to the join key (x.B or y.C) that contributed to this pair. Consequently, there is no way to tell how the new data batches ΔS_1 and ΔS_2 will contribute to the previous query results if we do not know how these results are related to the inputs S_1 and S_2. To compensate, we need to establish links between the query results and the data sources that were used to form their values. This is called *lineage tracking* and has been used for consistent representation of uncertain data [5] and for propagating annotations in relational queries [6]. In our case, this lineage tracking can be accomplished by propagating all keys used in joins and group-bys along with the values associated with the keys, so that, for each combination of keys, we have one group of result values. For our query, this is done by including the join key x.B as a group-by key. That is, in the following transformed query:

$$\begin{aligned} h(S_1, S_2) = \ & \textbf{select } ((\mathsf{x.A}, \mathsf{x.B}), (\textbf{sum}(\mathsf{y.D}), \textbf{count}(\mathsf{y.D}))) \\ & \textbf{from } \mathsf{x} \textbf{ in } S_1, \ \mathsf{y} \textbf{ in } S_2 \\ & \textbf{where } \mathsf{x.B} = \mathsf{y.C} \\ & \textbf{group by } \mathsf{x.A}, \mathsf{x.B} \end{aligned}$$

the join key is propagated to the output values so that the avg components, sum and count, are aggregations over groups that correspond to unique combinations of x.A and x.B. It can be shown that this query is a homomorphism

over S_1 and S_2, provided that the join is not on a many-to-many relationship. In general, a query with N joins/group-bys/order-bys will be transformed to a query that injects the join/group-by/order-by keys to the output so that each output value is annotated with a combination of N keys. Note that the output of the query h is larger than that of the original query q because it creates more groups and each group is assigned two values (sum and count), instead of one (avg). This is expected since we extended q with lineage tracking. The answer query a that gives the final result in $q(S_1, S_2) = a(h(S_1, S_2))$ is:

$$
\begin{array}{l}
a(X) = \textbf{select } (\mathsf{k},\ \textbf{sum}(\mathsf{s})/\textbf{sum}(\mathsf{c})) \\
\quad\quad \textbf{from } ((\mathsf{k},\mathsf{j}),(\mathsf{s},\mathsf{c})) \textbf{ in } X \\
\quad\quad \textbf{group by } \mathsf{k}
\end{array}
$$

that is, it removes the lineage j from X but also groups the result by the group-by key again and calculates the final avg values. The merge function for the homomorphism h is a full outer join on the lineage key θ that aggregates the matches. It is specified as:

$$
\begin{array}{l}
X \otimes Y = \textbf{select } (\theta,\ (\mathsf{sx}+\mathsf{sy},\mathsf{cx}+\mathsf{cy})) \\
\quad\quad \textbf{from } (\theta,(\mathsf{sx},\mathsf{cx})) \textbf{ in } X, \\
\quad\quad\quad\quad (\theta,(\mathsf{sy},\mathsf{cy})) \textbf{ in } Y \\
\quad\quad \textbf{union select } \mathsf{y} \textbf{ from } \mathsf{y} \textbf{ in } Y \\
\quad\quad\quad\quad \textbf{where } \pi_1(\mathsf{y}) \textbf{ not in } \pi_1(X) \\
\quad\quad \textbf{union select } \mathsf{x} \textbf{ from } \mathsf{x} \textbf{ in } X \\
\quad\quad\quad\quad \textbf{where } \pi_1(\mathsf{x}) \textbf{ not in } \pi_1(Y)
\end{array}
$$

where θ matches the lineage (k, j) and π_1 is pair and bag projection.

Merging States. The effectiveness of our incremental processing depends on the efficient implementation of the state transformation that merges the previous state with the results of processing the new data, state $\otimes\, h(\overline{\Delta S})$. The merge operation $X \otimes Y$ in the previous example can be implemented efficiently as a partitioned join. On Spark, for example, both the state and the new results are kept in the distributed memory of the worker nodes (as RDDs), while the full outer join can be implemented as a coGroup operation, which shuffles the join input data across the worker nodes using hash partitioning. However, when the new state is created by coGroup, it is already partitioned by the join key and is ready to be used for the next call to coGroup to handle the next batch of data. Consequently, only the results of processing the new data, which are typically smaller than the state, would have to be shuffled across the worker nodes before coGroup. Although other queries may require different merge functions, the correlation between the previous state and the results of processing the new data is always based on the lineage keys. Therefore, regardless of the query, we can keep the state partitioned on the lineage keys by simply leaving the new state partitions at the place they were generated. However, the new results would have to be partitioned and shuffled across the working nodes to be combined with the current state.

Our approach is based on the assumption that, since the state is kept in the distributed memory, there will be very little overhead in replacing the current state with a new state, as long as it is not repartitioned. But this assumption may not be valid if the input data or the current state is larger than the available distributed memory. Indeed, one of our goals is to be able to process data larger than the available memory, by processing these data incrementally, in batches that can fit in memory. However, if the state is stored on disk, replacing it with a new state may become prohibitively expensive. Fortunately, there is no need for using a distributed key-value store, such as HBase, to update the state in-place. There is, in fact, a simpler solution: since the state is kept partitioned on the lineage keys, every worker node may store its assigned state partition on its local disk, as a binary file (such as, an HDFS Sequence file) sorted by the lineage key. Then, after the results of processing the new data are distributed to worker nodes (using uniform hashing based on lineage key), each worker node will sort its new data partition by the lineage key and will merge it with its old state (stored on its local disk) to create a new state.

Iteration. Given that many important data analysis and mining algorithms, such as PageRank and k-means clustering, require repetition, we have extended our methods to include repetition, so that these algorithms too can become incremental. A repetition can take the following general form:

$$
\begin{array}{l}
X \leftarrow g(\overline{S}) \\
\textbf{for } i \leftarrow 1 \ldots n \\
\quad X \leftarrow f(X, \overline{S})
\end{array}
$$

where X is the fixpoint of the repetition that has initial value $g(\overline{S})$. That is, $X = f^n(g(\overline{S}), \overline{S})$, where f^n applies f n times. Note that X is not necessarily a bag. For example, non-negative Matrix Factorization [22], used in machine learning applications, such as for recommender systems, splits a matrix S into two non-negative matrices W and H. In that case, the fixpoint X is the pair (W, H), which is refined at every loop step.

An exact incremental solution of the above repetition is only possible if f is a homomorphism; a strict requirement that excludes many important iterative algorithms, such as PageRank and k-means clustering. Instead, our approach is to use an approximate solution that works well for iterative queries that improve a solution at each iteration step. As before, we split f into a homomorphism h, for some monoid $\otimes$, and an answer function a, so that $f(X, \overline{S}) = a(h(X, \overline{S}))$ and $h(X, \overline{S \uplus \Delta S}) = h(X, \overline{S}) \otimes h(X, \overline{\Delta S})$. Let T be the current state on input $\overline{S}$ and T' be the new state on input $\overline{S \uplus \Delta S}$. An ideal solution would have been to derive the new state T' incrementally, as $T \otimes \Delta T$, so that the state increment ΔT depends on ΔS but not on T. But this independence from T is not always possible for many iterative algorithms. In PageRank, for example, we cannot just calculate the PageRanks of the new data and merge them with the PageRank of the existing data because the new PageRank contributions may have to propagate to the rest of the graph. On the other hand, it would be too expensive to correlate the entire state T with ΔT at each iteration step. Our compromise is to partially correlate T with ΔT at each iteration step, and then fully merge ΔT with T after the iteration. This partial correlation is accomplished with the operation $\widehat{\otimes}$, called *diffusion*, which is directly derived from the merge function $\otimes$. That is, while $T \otimes \Delta T$ returns a new complete state by merging T with ΔT, the operation $T \,\widehat{\otimes}\, \Delta T$ returns a new ΔT that contains only the part of $T \otimes \Delta T$ that is different from T. Although $T \,\widehat{\otimes}\, \Delta T$ is larger than ΔT, it is often far smaller than $T \otimes \Delta T$. For instance, if T were a bag,

$T \widehat{\otimes} \Delta T$ would have been a subset of $T \otimes \Delta T$. Based on this analysis, we are using the following approximate algorithm to compute the iteration incrementally:

$$
\begin{array}{l}
\Delta X \leftarrow g(\overline{\Delta S}) \\
\textbf{for } i \leftarrow 1 \ldots n \\
\quad \Delta T \leftarrow T \widehat{\otimes} h(\Delta X, \overline{\Delta S}) \\
\quad \Delta X \leftarrow a(\Delta T) \\
T \leftarrow T \otimes \Delta T \\
\textbf{return } a(T)
\end{array}
$$

The diffusion operator $\widehat{\otimes}$ must satisfy the property:

$$T \otimes (T \widehat{\otimes} \Delta T) = T \otimes (T \otimes \Delta T)$$

that is, when $T \widehat{\otimes} \Delta T$ and $T \otimes \Delta T$ are merged with T, they give the same answer, because $T \widehat{\otimes} \Delta T$ discards the parts of T that are not joined with ΔT, but these parts are embedded back in the answer when we merge with T.

For example, PageRank is an iterative algorithm that calculates the PageRank of a graph node as the sum of the incoming PageRank contributions from its neighbors, while its own PageRank is equally distributed to its outgoing neighbors. The PageRank query expressed in MRQL is a simple self-join on the graph, which is optimized into a single group-by operation [18]. For PageRank, the state merging $\otimes$ is a full outer join that incorporates the new PageRank contributions to the existing PageRanks. The diffusion operation $T \widehat{\otimes} \Delta T$, on the other hand, propagates the new PageRank contributions from ΔT to T, that is, from the nodes in ΔT to their immediate outgoing neighbors in both T and ΔT. Thus, at each iteration step, ΔT is expanded to $T \widehat{\otimes} \Delta T$, growing one level at a time, in a way similar to breadth-first-search. Consequently, our approximate algorithm propagates PageRanks up to depth n, starting from the ΔT nodes, and only the affected nodes will be part of the new ΔT. The $\widehat{\otimes}$ operation is similar to $\otimes$, but with a right-outer join instead of a full outer join. That way, the data from T that are not joined with ΔT will not appear in the new ΔT.

A more complete example is the following query that implements the k-means clustering algorithm by repeatedly deriving k new centroids from the old:

```
repeat centroids = ...
  step select < X: avg(s.X), Y: avg(s.Y) >
         from s in S1
       group by k: (select c from c in centroids
                    order by distance(c,s))[0]
```

where S_1 is the input stream of points on the X-Y plane, centroids is the current set of centroids (k cluster centers), and distance is a function that calculates the distance between two points. The initial value of centroids (the ... value) can be a bag of k random points. The inner select-query in the group-by part assigns the closest centroid to a point s (where [0] returns the first tuple of an ordered list). The outer select-query in the repeat step clusters the data points by their closest centroid, and, for each cluster, a new centroid is calculated from the average values of its points. As in the previous join-groupBy example, the average value of a bag of values is decomposed into a pair that contains the sum and the count of values. That is, the state is a bag of $\{(k, ((sx, cx), (sy, cy)))\}$ so that the centroids are the bag of points $\{(sx/cx, sy/cy)\}$ and k is the lineage (the group-by key), which is the X-Y coordinates of a centroid. Consequently, the answer query that returns the final result (the centroids) is:

```
a(state) = select < X: sx/cx, Y: sy/cy >
           from (k,((sx,cx),(sy,cy))) in state
```

(it does not require a group-by since k is the only lineage key), while $h(X, S_1)$ is:

```
select (k, ( (sum(s.X),count(s.X)),
             (sum(s.Y),count(s.Y)) ))
  from s in S1
 group by k: (select c from c in X
              order by distance(c,s))[0]
```

The merge function $\otimes$ is a full outer join, similar to the one used by the join-groupBy example. The diffusion operation $X \widehat{\otimes} Y$ though is a right-outer join that discards those state data that do not join with the new data. That is, it is equal to the $X \otimes Y$ query without the last union:

```
X ⊗̂ Y = select (k, (sx+sy,cx+cy))
         from (k,(sx,cx)) in X,
              (k,(sy,cy)) in Y
       union select y from y in Y
             where π1(y) not in π1(X)
```

But, suppose that one of the group-by or join keys in a query is a floating point number. This is the case with the previous k-means clustering query, because it groups the points by their closest centroids, which contain floating point numbers. The group-by operation itself is not a problem because the centroids are fixed during the group-by. The problem arises when merging the current state with the results of processing the new data. For the k-means query example, the merge function $\otimes$ is an equi-join whose join attribute is a centroid, so that the sum and count values associated with the same centroid are brought together from the join inputs and are accumulated. Since the join condition is over attributes with floating point numbers, the join condition will fail in most cases. This problem becomes even worse for the approximate solution, because it uses different sets of centroids X and X_{prev} when states are merged. Most iterative queries do not have this problem. The lineage in PageRank, for example, is the node ID, which remains invariant across iterations. For these uncommon queries, such as k-means, that have floating point numbers in their join/group-by/order-by attributes, we use yet another approximation: the join is done based on an "approximate equality" where two floating point numbers are taken to be equal if their difference is below some given threshold. This works well for our approximate solution for iteration because it is based on the assumption that the new solution X is approximately equal to the previous one, X_{prev}.

3. RELATED WORK

New frameworks in distributed Big Data analytics have become essential tools to large-scale machine learning and scientific discoveries. Among these frameworks, the Map-Reduce programming model [14] has emerged as a generic, scalable, and cost effective solution for Big Data processing on clusters of commodity hardware. One of the major drawbacks of the Map-Reduce model is that, to simplify reliability and fault tolerance, it does not preserve data

in memory between the map and reduce tasks of a Map-Reduce job or across consecutive jobs, which imposes a high overhead to complex workflows and graph algorithms, such as PageRank, which require repetitive Map-Reduce jobs. Recent systems for cloud computing use distributed memory for inter-node communication, such as the main memory Map-Reduce (M3R [38]), Apache Spark [40], Apache Flink [20], Piccolo [36], and distributed GraphLab [28]. Another alternative framework to the Map-Reduce model is the Bulk Synchronous Parallelism (BSP) programming model [44]. The best known implementations of the BSP model for data analysis on the cloud are Google's Pregel [29], Apache Giraph [21], and Apache Hama [25].

Although the Map-Reduce framework was originally designed for batch processing, there are several recent systems that have extended Map-Reduce with online processing capabilities. Some of these systems build on the well-established research on data streaming based on sliding windows and incremental operators [4], which includes systems such as Aurora [1] and Telegraph [11]. MapReduce Online [12] maintains state in memory for a chain of MapReduce jobs and reacts efficiently to additional input records. It also provides a memoization-aware scheduler to reduce communication across a cluster. Incoop [7] is a Hadoop-based incremental processing system with an incremental storage system that identifies the similarities between the input data of consecutive job runs and splits the input based on the similarity and file content. i^2MapReduce [47] implements incremental iterative Map-Reduce jobs using a store, MRB-Store, that maps input values to the reduce output values. This store is used for detecting delta changes and propagating these changes to the output. Google's Percolator [35] is a system based on BigTable for incrementally processing updates to a large data set. It updates an index incrementally as new documents are crawled. Microsoft Naiad [32] is a distributed framework for cyclic dataflow programs that facilitates iterative and incremental computations. It is based on differential dataflow computations, which allow incremental computations to have nested iterations. CBP [27] is a continuous bulk processing system on Hadoop that provides a stateful group-wise operator that allows users to easily store and retrieve state during the reduce stage as new data inputs arrive. Their incremental computing PageRank implementation is able to cut running time in half. REX [31] handles iterative computations in which changes in the form of deltas are propagated across iterations and state is updated efficiently. In contrast to our automated approach, REX requires the programmer to explicitly specify how to process deltas, which are handled as first class objects. Trill [10] is a high throughput, low latency streaming query processor for temporal relational data, developed at Microsoft Research. The Reactive Aggregator [43], developed at IBM Research, is a new sliding-window streaming engine that performs many forms of sliding-window aggregation incrementally. In addition to these general data analysis engines, there are many data analysis algorithms that have been implemented incrementally, such as incremental pagerank [15]. Finally, the incremental query processing is related to the problem of incremental view maintenance, which has been extensively studied in the context of relational views (see [23] for a literature survey).

Many novel Big Data stream processing systems, also known as distributed stream processing engines (DSPEs), have emerged recently. The most popular one is Twitter's Storm [35], which is now part of the Apache ecosystem for Big Data analytics. It provides primitives for transforming streams based on a user-defined topology, consisting of spouts (stream sources) and bolts (which consume input streams and may emit new streams). Other popular DSPE platforms include Spark's D-Streams [46], Flink Streaming [20], Apache S4 [42], and Apache Samza [39].

In programming languages, self-adjusting computation [2] refers to a technique for compiling batch programs into programs that can automatically respond to changes to their data. It requires the construction of a dependence graph at run-time so that when the computation data changes, the output can be updated by re-evaluating only the affected parts of the computation. In contrast to our work, which requires only the state to reside in memory, self-adjusting computation expects both the input and the output of a computation to reside in memory, which makes it inappropriate for unbounded data in a continuous stream. Furthermore, such dynamic methods impose a run-time storage and computation overhead by maintaining the dependence graph. The main idea in [2] is to manually annotate the parts of the input type that is changeable, and the system will derive an incremental program automatically based on these annotations. Each changeable value is wrapped by a mutator that includes a list of reader closures that need to be evaluated when the value changes. A read operation on a mutator inserts a new closure, while the write operation triggers the evaluation of the closures, which may cause writes to other mutators, etc, resulting to a cascade of closure execution triggered by changed data only. This technique has been extended to handle incremental list insertions (like our work), but it requires the rewriting of all list operations to work on incremental lists. Recently, there is a proof-of-concept implementation of this technique on map-reduce [3], but it was tested on a serial machine. It is doubtful that such dynamic techniques can be efficiently applied to a distributed environment, where a write in one compute node may cause a read in another node. Finally, there is recent work on static incrementalization based on derivatives [8]. In contrast to our work, it assumes that the merge function that combines the previous result with the result on the delta changes uses exactly the same delta changes, a restriction that excludes aggregations and group-bys.

4. EARLIER WORK: MRQL

Apache MRQL [33] is a query processing and optimization system for large-scale, distributed data analysis. MRQL was originally developed by the author ([18, 17]), but is now an Apache incubating project with many developers and users worldwide. The MRQL language is an SQL-like query language for large-scale data analysis on computer clusters. The MRQL query processing system can evaluate MRQL queries in four modes: in Map-Reduce mode using Apache Hadoop [24], in BSP mode (Bulk Synchronous Parallel model) using Apache Hama [25], in Spark mode using Apache Spark [40], and in Flink mode using Apache Flink [20]. The MRQL query language is powerful enough to express most common data analysis tasks over many forms of raw in-situ data, such as XML and JSON documents, binary files, and CSV documents. The design of MRQL has been influenced by XQuery and ODMG OQL, although

it uses SQL-like syntax. In fact, when restricted to XML, MRQL is as powerful as XQuery. MRQL is more powerful than other current high-level Map-Reduce languages, such as Hive [26] and PigLatin [34], since it can operate on more complex data and supports more powerful query constructs, thus eliminating the need for using explicit procedural code. With MRQL, users are able to express complex data analysis tasks, such as PageRank, k-means clustering, matrix factorization, etc, using SQL-like queries exclusively, while the MRQL query processing system is able to compile these queries to efficient Java code that can run on various distributed processing platforms. For example, the PageRank query on raw DBLP XML data, which ranks authors based on the number of citations they have received from other authors, is 16 lines long [17] and can be executed on all the supported platforms as is, without changing the query.

A recent extension to MRQL, called *MRQL Streaming*, supports the processing of continuous MRQL queries over streams of batch data (that is, data that come in continuous large batches). Before the incremental MRQL work presented in this paper, MRQL Streaming supported traditional window-based streaming based on a fixed window during a specified time interval. Any batch MRQL query can be converted to a window-based streaming query by replacing at least one of the 'source' calls in the query that access data sources to 'stream' calls (with exactly the same call arguments). For example, the query:

```
select (k,avg(p.Y))
from p in stream(binary,"points")
group by k: p.X
```

groups a stream of points by their X coordinate and returns the average Y values in each group. The MRQL Streaming engine first processes all the existing sequence files in the directory points and then checks this directory periodically for new files. When new files are inserted in the directory, it processes the new batch of data using distributed query processing. MRQL Streaming also supports a stream input format for listening to TCP sockets for text input based on one of the MRQL Parsed Input Formats (XML, JSON, CSV). A query may work on multiple stream sources and multiple batch sources. If there is at least one stream source, the query becomes continuous (it never stops). The output of a continuous query is stored in a file directory, where each file contains the results of processing each batch of streaming data. Currently, MRQL Streaming works on Spark Streaming only but there are current efforts to add support for Storm and Flink Streaming in the near future. The work reported here, called *Incremental MRQL*, extends the current MRQL Streaming engine with incremental stream processing.

5. THE MRQL ALGEBRA

Our compiler translates queries to algebraic terms and then uses rewrite rules to put these algebraic terms into a homomorphic form, which is then used to compute the query results incrementally by combining the previous results with the results of processing the incremental batches.

The MRQL algebra described in this section is a variation of the algebra presented in our previous work [18], but is more suitable for describing our incremental methods. The relational algebra, the nested relational algebra, as well as many other database algebras can be easily translated to our algebra. Our algebra consists of a small number of higher-order homomorphic operators [18], which are defined using structural recursion based on the union representation of bags [19]. Monoid homomorphisms capture the essence of many divide-and-conquer algorithms and can be used as the basis for data parallelism [19].

The first operator, cMap (also known as concat-map or flatten-map in functional programming languages), generalizes the select, project, join, and unnest operators of the nested relational algebra. Given two arbitrary types α and β, the operation $\mathrm{cMap}(f, X)$ maps a bag X of type $\{\alpha\}$ to a bag of type $\{\beta\}$ by applying the function f of type $\alpha \to \{\beta\}$ to each element of X, yielding one bag for each element, and then by merging these bags to form a single bag of type $\{\beta\}$. Using a set former notation on bags, it is expressed as:

$$\mathrm{cMap}(f, X) \triangleq \{\, z \,|\, x \in X,\, z \in f(x) \,\} \tag{1}$$

or, alternatively, using structural recursion:

$$\begin{array}{rcl} \mathrm{cMap}(f, X \uplus Y) & = & \mathrm{cMap}(f, X) \uplus \mathrm{cMap}(f, Y) \\ \mathrm{cMap}(f, \{a\}) & = & f(a) \\ \mathrm{cMap}(f, \{\,\}) & = & \{\,\} \end{array}$$

Given an arbitrary type κ that supports value equality ($=$), an arbitrary type α, and a bag X of type $\{(\kappa, \alpha)\}$, the operation $\mathrm{groupBy}(X)$ groups the elements of the bag X by their first component and returns a bag of type $\{(\kappa, \{\alpha\})\}$. For example, groupBy({ (1,"A"), (2,"B"), (1,"C") }) returns {(1,{"A","C"}), (2,{"B"})}. The groupBy operation cannot be defined using a set former notation, but can be defined using structural recursion:

$$\begin{array}{rcl} \mathrm{groupBy}(X \uplus Y) & = & \mathrm{groupBy}(X) \Uparrow_{\uplus} \mathrm{groupBy}(Y) \\ \mathrm{groupBy}(\{(k, a)\}) & = & \{(k, \{a\})\} \\ \mathrm{groupBy}(\{\,\}) & = & \{\,\} \end{array}$$

where the parametric monoid $\Uparrow_{\oplus}$ is a full outer join that merges groups associated with the same key using the monoid $\oplus$ (equal to $\uplus$ for groupBy):

$$\begin{array}{rcl} X \Uparrow_{\oplus} Y & \triangleq & \{\, (k, a \oplus b) \,|\, (k, a) \in X,\, (k, b) \in Y \,\} \\ & & \uplus\, \{\, (k, a) \,|\, (k, a) \in X,\, k \notin \pi_1(Y) \,\} \\ & & \uplus\, \{\, (k, b) \,|\, (k, b) \in Y,\, k \notin \pi_1(X) \,\} \end{array} \tag{2}$$

where $\pi_1(X) = \{\, k \,|\, (k, x) \in X \,\}$. In other words, the monoid $\Uparrow_{\uplus}$ constructs a set of pairs whose unique key is the first pair element. In fact, any bag X can be converted to a set using $\pi_1(\mathrm{groupBy}(\mathrm{cMap}(\lambda x.\, \{(x, x)\},\, X)))$. Note also that $\mathrm{groupBy}(X)$ is not the same as the nesting $\{\, (k, \{\, y \,|\, (k', y) \in X,\, k = k' \,\}) \,|\, (k, x) \in X \,\}$, as the latter contains duplicate entries for the key k. Unlike nesting, unnesting a groupBy returns the input bag (proven in Appendix A):

$$\{\, (k, v) \,|\, (k, s) \in \mathrm{groupBy}(X),\, v \in s \,\} \;=\; X \tag{3}$$

Although any join $X \bowtie_p Y$ can be expressed as a nested cMap:

$$\mathrm{cMap}(\lambda x.\, \mathrm{cMap}(\lambda y.\, \mathrm{if}\ p(x, y)\ \mathrm{then}\ \{(x, y)\}\ \mathrm{else}\ \{\,\},\, Y),\, X)$$

this term is not always a homomorphism on both inputs. Instead, MRQL provides a special homomorphic operation for equi-joins and outer joins, $\mathrm{coGroup}(X, Y)$, between a bag X of type $\{(\kappa, \alpha)\}$ and a bag Y of type $\{(\kappa, \beta)\}$ over their first component of a type κ, which returns a bag of

type $\{(\kappa, (\{\alpha\}, \{\beta\}))\}$:

$$\begin{aligned}
\text{coGroup}(X_1 \uplus X_2, Y_1 \uplus Y_2) &= \text{coGroup}(X_1, Y_1) \\
&\quad \Uparrow_{\uplus \times \uplus} \text{coGroup}(X_2, Y_2) \\
\text{coGroup}(\{(k,a)\}, \{(k,b)\}) &= \{(k, (\{a\}, \{b\}))\} \\
\text{coGroup}(\{(k,a)\}, \{(k',b)\}) &= \{(k, (\{a\}, \{\,\})), \\
&\quad (k', (\{\,\}, \{b\}))\} \\
\text{coGroup}(\{(k,a)\}, \{\,\}) &= \{(k, (\{a\}, \{\,\}))\} \\
\text{coGroup}(\{\,\}, \{(k,b)\}) &= \{(k, (\{\,\}, \{b\}))\} \\
\text{coGroup}(\{\,\}, \{\,\}) &= \{(k, (\{\,\}, \{\,\}))\}
\end{aligned}$$

where the product of two monoids, $\oplus \times \otimes$ is a monoid that, when applied to two pairs (x_1, x_2) and (y_1, y_2), returns $(x_1 \oplus y_1, x_2 \otimes y_2)$. That is, the monoid $\Uparrow_{\uplus \times \uplus}$ merges two bags of type $\{(\kappa, (\{\alpha\}, \{\beta\}))\}$ by unioning together their $\{\alpha\}$ and $\{\beta\}$ values that correspond to the same key κ. For example, coGroup({ (1,"A"), (2,"B"), (1,"C") }, { (1,"D"), (2,"E"), (3,"F") }) returns {(1,({"A","C"},{"D"})), (2,({"B"},{"E"})), (3,({},{"F"}))}. It can be proven (with a proof similar to that of Equation (3)) that both coGroup inputs can be derived from the coGroup result:

$$\begin{aligned}
\{\,(k,x) \,|\, (k,(s_1,s_2)) \in \text{coGroup}(X,Y),\, x \in s_1\,\} &= X \\
\{\,(k,y) \,|\, (k,(s_1,s_2)) \in \text{coGroup}(X,Y),\, y \in s_2\,\} &= Y
\end{aligned}$$

and a coGroup is equivalent to a groupBy if one of the inputs is empty:

$$\text{coGroup}(X, \{\,\}) = \{\,(k, (s, \{\,\})) \,|\, (k,s) \in \text{groupBy}(X)\,\}$$

Aggregations are captured by the operation $\text{reduce}(\oplus, X)$, which aggregates a bag X using a commutative monoid $\oplus$. For example, $\text{reduce}(+, \{1,2,3\}) = 6$. This operation is in fact a general homomorphism that can be defined on any monoid $\oplus$ with an identity function. We have:

$$\begin{aligned}
\text{reduce}(\uplus, X) &= X \\
\text{reduce}(\Uparrow_{\uplus}, X) &= \text{groupBy}(X)
\end{aligned}$$

Finally, iteration $\text{repeat}(F, n, X)$ over the bag X of type $\{\alpha\}$ applies F of type $\{\alpha\} \to \{\alpha\}$ to X n times, yielding a bag of type $\{\alpha\}$:

$$\text{repeat}(F, n, X) = F^n(X) = F(F(\ldots F(X)))$$

An iteration is a homomorphism as long as F is a homomorphism, that is, when $F(X \uplus Y) = F(X) \uplus F(Y)$.

Definition 1 (MRQL Algebra)**.** *The MRQL algebra consists of terms that take the following form:*

e, e_1, e_2	::=	$\text{cMap}(f, e)$	*flatten-map*
	|	$\text{groupBy}(e)$	*group-by*
	|	$\text{coGroup}(e_1, e_2)$	*join*
	|	$\text{reduce}(\oplus, e)$	*aggregation*
	|	S_i	*stream source*

where $\oplus$ is a monoid on basic types, such as $+$, $*$, or, and, etc. Function f is an anonymous function that may contain such algebraic terms but is not permitted to contain any reference to a stream source, S_i.

The restriction on f in Definition 1 excludes non-equi-joins, such as cross products, which require a nested cMap in which the inner cMap is over a data source. This algebra does not include iteration, $\text{repeat}(F, n, e)$. Iterations have been discussed in Section 2. In addition, for brevity, this algebra does not include terms for non-streaming input sources, general tuple and record construction and projection, bag union, singleton and empty bag, arithmetic operations, if-then-else expressions, boolean operations, etc.

In addition to these operations, there are a few more algebraic operations that can be expressed as homomorphisms, such as $\text{orderBy}(X)$, which is a groupBy followed by a sorting over the group-by key. This operation returns a list, which is represented by the non-commutative monoid list-append, $+\!\!+$. Mixing multiple collection monoids and operations in the same algebra has been addressed by our previous work [19], and is left out from this paper to keep our analysis simple.

For example, the query $q(S_1, S_2)$ used as the first example in Section 2, is translated to the following algebraic term $Q(S_1, S_2)$:

$$\begin{aligned}
&\text{cMap}(\lambda(k,s).\, \{(k, avg(s))\}, \\
&\quad \text{groupBy}(\text{cMap}(\lambda(j, (xs, ys)).\, g(xs, ys), \\
&\qquad \text{coGroup}(\text{cMap}(\lambda x.\, \{(x.B, x)\},\, S_1), \\
&\qquad\quad \text{cMap}(\lambda y.\, \{(y.C, y)\},\, S_2)))))
\end{aligned}$$

where $avg(s) = \text{reduce}(+, s)/\text{reduce}(+, \text{cMap}(\lambda v.\, \{1\}, s))$ and $g(xs, ys) = \text{cMap}(\lambda x.\, \text{cMap}(\lambda y.\, \{(x.A, y.D)\},\, ys),\, xs)$.

Although all algebraic operators used in MRQL are homomorphisms, their composition may not be. For instance, $\text{cMap}(f, \text{groupBy}(X))$ is not a homomorphism for certain functions f, because, in general, cMap does not distribute over $\Uparrow_{\uplus}$. One of our goals is to transform any composition of algebraic operations into a homomorphism.

6. QUERY NORMALIZATION

In an earlier work [18], we have presented general algorithms for unnesting nested queries. For example, consider the following nested query:

```
select x from x in X
where x.D > sum(select y.C from y in Y where x.A=y.B)
```

A typical method for evaluating this query in a relational system is to first group Y by y.B, yielding pairs of y.B and sum(y.C), and then to join the result with X on x.A=y.B using a left-outer join, removing all those matches whose x.D is below the sum. But, in our framework, this query is translated into:

```
cMap( λ(k,(xs,ys)). cMap( λx. if x.D > reduce(+,ys)
                              then {x}
                              else {},  xs),
      coGroup( cMap( λx. {(x.A,x)}, X ),
               cMap( λy. {(y.B,y.C)}, Y ) ) )
```

That is, the query unnesting is done with a left-outer join, which is captured concisely by the coGroup operation without the need for using an additional group-by operation or handling null values. This unnesting technique was generalized to handle arbitrary nested queries, at any place, number, and nesting level (the reader is referred to our earlier work [18] for details).

The algebraic terms derived from MRQL queries can be normalized using the following rule:

$$\text{cMap}(f, \text{cMap}(g, S)) \to \text{cMap}(\lambda x.\, \text{cMap}(f, g(x)), S) \quad (4)$$

which fuses two cascaded cMaps into a nested cMap, thus avoiding the construction of the intermediate bag. This rule

can be proven directly from the cMap definition in Equation (1):

$$\begin{array}{l}\text{cMap}(f, \text{cMap}(g, S)) \\ = \{\, z \,|\, w \in \{\, y \,|\, x \in S,\, y \in g(x) \,\},\, z \in f(w) \,\} \\ = \{\, z \,|\, x \in S,\, y \in g(x),\, z \in f(y) \,\} \\ = \{\, z \,|\, x \in S,\, z \in \{\, w \,|\, y \in g(x),\, w \in f(y) \,\} \,\} \\ = \text{cMap}(\lambda x.\, \text{cMap}(f, g(x)), S)\end{array}$$

If we apply the transformation (4) repeatedly, and given that we can always use the identity $\text{cMap}(\lambda x.\{x\}, X) = X$ in places where there is no cMap between groupBy/coGroup operations, any algebraic terms in Definition 1 can be normalized into the following form:

Definition 2 (Normalized MRQL Algebra)**.** *The normalized MRQL algebra consists of terms q that take the following form:*

$$\begin{array}{rcll} q, q_1, q_2 & ::= & \text{reduce}(\oplus,\, c) & \textit{the query header} \\ & | & c & \\ c, c_1, c_2 & ::= & \text{cMap}(f,\, e) & \\ e & ::= & \text{groupBy}(c) & \textit{the query body} \\ & | & \text{coGroup}(c_1,\, c_2) & \\ & | & S_i & \end{array}$$

where function f is an anonymous function that does not contain any reference to a stream source, S_i.

The query body is a tree of groupBy/coGroup operations connected via cMaps.

7. MONOID INFERENCE

One of our tasks is, given an algebraic term $f(\overline{S})$, where an $S_i \in \overline{S}$ is a streaming data source, to prove that f is a homomorphism by deriving a monoid $\otimes$ such that:

$$f(S_1 \uplus S_1', \ldots, S_n \uplus S_n') = f(S_1, \ldots, S_n) \otimes f(S_1', \ldots, S_n') \quad (6)$$

We have developed a monoid inference system, inspired by type inference systems used in programming languages. We use the judgment $\rho \vdash e : \oplus$ to indicate that e is a monoid homomorphism with a merge function $\oplus$ under the environment ρ, which binds variables to monoids. The notation $\rho(v)$ extracts the binding of the variable v, while $\rho[v : \oplus]$ extends the environment with a new binding from v to $\oplus$. Equation (6) can now be expressed as the judgment:

$$[S_1 : \uplus, \ldots, S_n : \uplus] \vdash f(\overline{S}) : \otimes$$

If a term is invariant under change, such as an invariant data source, it is associated with the special monoid $\Box$:

$$X \,\Box\, Y \;\triangleq\; \begin{cases} X & \text{if } X = Y \\ \text{error} & \text{otherwise} \end{cases}$$

Our monoid inference algorithm is a heuristic algorithm that annotates terms with monoids (when possible). It is very similar to type inference. Most of our inference rules are expressed as fractions: the denominator (below the line) contains the premises (separated by comma) and the numerator (above the line) is the conclusion. For example, the rule $\frac{\rho \vdash f(x) : \otimes}{\rho \vdash x : \oplus}$ indicates that $f(x_1 \oplus x_2) = f(x_1) \otimes f(x_2)$. Figure 2 gives some of the inference rules. More rules will be given in Lemma 1. Rules (5c) through (5f) are derived directly from the algebraic definition of the operators. Rule (5a) retrieves the associated monoid of a variable v from the environment ρ. Rule (5b) indicates that if all the variables in a term e are invariant, then so is e. Rule (5h) indicates that a cMap over a groupBy is a homomorphism as long as its functional argument is a homomorphism. It can be proven as follows:

$$\begin{array}{l}\text{cMap}(f,\, X \Uparrow_\oplus Y) \\ = \{\, (\theta, z) \,|\, (k, s) \in (X \Uparrow_\oplus Y),\, (\theta, z) \in f(k, s) \,\} \\ = \{\, (\theta, z) \,|\, (k, x) \in X,\, (k, y) \in Y,\, (\theta, z) \in f(k, x \oplus y) \,\} \uplus \cdots \\ = \{\, (\theta, z) \,|\, (k, x) \in X,\, (k, y) \in Y, \\ \qquad (\theta, z) \in (f(k, x) \Uparrow_\otimes f(k, y)) \,\} \uplus \cdots \\ = \{\, (\theta, z) \,|\, (k, x) \in X,\, (k, y) \in Y, \\ \qquad (\theta, z) \in \{\, (\theta, a \otimes b) \,|\, (\theta, a) \in f(k, x), \\ \qquad\qquad (\theta, b) \in f(k, y) \,\} \,\} \uplus \cdots \\ = \{\, (\theta, a \otimes b) \,|\, (k, x) \in X,\, (k, y) \in Y,\, (\theta, a) \in f(k, x), \\ \qquad (\theta, b) \in f(k, y) \,\} \uplus \cdots \\ = \{\, (\theta, a \otimes b) \,|\, (\theta, a) \in \text{cMap}(f, X),\, (\theta, b) \in \text{cMap}(f, Y) \,\} \\ \quad \uplus \cdots \qquad \text{(given that } k \text{ and } \theta \text{ are unique keys)} \\ = \text{cMap}(f, X) \;\Uparrow_\otimes\; \text{cMap}(f, Y)\end{array}$$

8. INJECTING LINEAGE TRACKING

There are two tasks that need to be accomplished to achieve our goal of transforming an algebraic term into a homomorphism: 1) transform the algebraic term in such a way that it propagates all keys used in joins and group-bys to the query output, and 2) pull the non-homomorphic parts out of the algebraic term so that it becomes a homomorphism. In this section, we address the first task.

We will transform the algebraic terms given in Definition 2 in such a way that they propagate the join and the group-by keys. That is, each value v returned by these terms is annotated with a lineage θ, as a pair (θ, v), where θ takes the following form:

$$\begin{array}{rcll} \theta, \theta_1, \theta_2 & ::= & (\kappa, \theta) & \text{extended with groupBy key } \kappa \\ & | & (\kappa, (\theta_1, \theta_2)) & \text{extended with coGroup key } \kappa \\ & | & () & \text{empty lineage} \end{array}$$

That is, the lineage θ of the query result v is the tree of the groupBy and coGroup keys that are used in deriving the result v (one key for each groupBy and coGroup operation). The lineage tree θ has the same shape as the groupBy/coGroup tree of the query. We transform a query q in such a way that, if the output of the query is $\{t\}$ for some type t, then the transformed query will have output $\{(\theta, t)\}$. Furthermore, if the the output is a non-collection type t, then the transformed query will also have output $\{(\theta, t)\}$, which separates the contributions to t associated with each combination of group-by/join keys.

Algorithm 1 (Lineage Annotation)**.**
Input: *a normalized query term q defined in Definition 2*
Output: *a term $\mathcal{T}_q[\![q]\!]$ annotated with a lineage θ*

$$\rho \vdash v : \rho(v) \quad (5a)$$

$$\frac{\rho \vdash e : \square}{\forall v \in e : \ \rho \vdash v : \square} \quad (5b)$$

$$\frac{\rho \vdash \text{reduce}(\oplus, X) : \oplus}{\rho \vdash X : \uplus} \quad (5c)$$

$$\frac{\rho \vdash (X \uplus Y) : \uplus}{\rho \vdash X : \uplus, \ \rho \vdash Y : \uplus} \quad (5d)$$

$$\frac{\rho \vdash \text{cMap}(f, X) : \uplus}{\rho \vdash X : \uplus} \quad (5e)$$

$$\frac{\rho \vdash \text{groupBy}(X) : \Uparrow_{\uplus}}{\rho \vdash X : \uplus} \quad (5f)$$

$$\frac{\rho \vdash \text{coGroup}(X, Y) : \Uparrow_{\uplus \times \uplus}}{\rho \vdash X : \uplus, \ \rho \vdash Y : \uplus} \quad (5g)$$

$$\frac{\rho \vdash \text{cMap}(f, X) : \Uparrow_{\otimes}}{\rho \vdash X : \Uparrow_{\oplus}, \ \rho[k : \square, s : \oplus] \vdash f(k, s) : \Uparrow_{\otimes}} \quad (5h)$$

Figure 2: Monoid Inference Rules

$$\mathcal{T}_q[\![\text{reduce}(\oplus, \text{cMap}(f, S_i))]\!] = \{((), \text{reduce}(\oplus, \text{cMap}(f, S_i)))\} \quad (7a)$$

$$\mathcal{T}_q[\![\text{reduce}(\oplus, \text{cMap}(f, e))]\!] = \text{reduce}(\Uparrow_{\oplus}, \text{sMap1}(f, \mathcal{T}_e[\![e]\!])) \quad (7b)$$

$$\mathcal{T}_q[\![\text{cMap}(f, S_i)]\!] = \{ ((), b) \mid a \in S_i, \ b \in f(a) \} \quad (7c)$$

$$\mathcal{T}_q[\![\text{cMap}(f, e)]\!] = \text{sMap1}(f, \mathcal{T}_e[\![e]\!]) \quad (7d)$$

$$\mathcal{T}_e[\![\text{groupBy}(c)]\!] = \text{groupBy}(\text{swap}(\mathcal{T}_c[\![c]\!])) \quad (7e)$$

$$\mathcal{T}_e[\![\text{coGroup}(c_1, c_2)]\!] = \text{mix}(\text{coGroup}(\mathcal{T}_c[\![c_1]\!], \mathcal{T}_c[\![c_2]\!])) \quad (7f)$$

$$\mathcal{T}_c[\![\text{cMap}(f, S_i)]\!] = \text{sMap3}(f, S_i) \quad (7g)$$

$$\mathcal{T}_c[\![\text{cMap}(f, e)]\!] = \text{sMap2}(f, \mathcal{T}_e[\![e]\!]) \quad (7h)$$

where sMap1, sMap2, sMap3, swap, and mix are defined as follows:

$$\text{sMap1}(f, X) \triangleq \{ ((\theta, k), b) \mid ((\theta, k), a) \in X, \ b \in f(k, a) \} \quad (8a)$$

$$\text{sMap2}(f, X) \triangleq \{ (k', ((k, \theta), b)) \mid ((k, \theta), a) \in X, \ (k', b) \in f(k, a) \} \quad (8b)$$

$$\text{sMap3}(f, X) \triangleq \{ (k, ((), b)) \mid a \in X, \ (k, b) \in f(a) \} \quad (8c)$$

$$\text{swap}(X) \triangleq \{ ((k, \theta), v) \mid (k, (\theta, v)) \in X \} \quad (8d)$$

$$\text{mix}(X) \triangleq \{ ((k, (\theta_x, \theta_y)), (xs, ys)) \mid (k, (s_1, s_2)) \in X, \ (\theta_x, xs) \in \text{groupBy}(s_1), \ (\theta_y, ys) \in \text{groupBy}(s_2) \} \quad (8e)$$

A query q in our framework is transformed in such a way that it propagates the lineage from the data stream sources to the query output, starting with the empty lineage () at the sources and extended with the join and group-by keys. The sMap1 operation is a cMap that propagates the input lineage θ to the output as is. The sMap2 operation is a cMap that extends the input lineage θ with a groupBy/coGroup key k. The lineage propagation is done by the cMap Rules (7c) and (7h). Rule (7c) applies to the outer query cMap that produces the query output. It simply propagates the lineage from the cMap input to the output. Rule (7h) applies to a cMap that returns the input of a groupBy or coGroup. The output of this cMap must be a bag of key-value pairs, as is expected by a groupBy or a coGroup. Thus, Rule (7h) extends the lineage with a new key and prepares the cMap output for the enclosing groupBy or coGroup. This is done by translating cMap to sMap2. Rule (7b) indicates that a total aggregation becomes a group-by aggregation by aggregating the values of each group associated with a different lineage θ. Rule (7e) translates a groupBy on a key to a groupBy on the entire lineage (which includes the groupBy key). Rule (7f) translates a coGroup on a key to a coGroup on the entire lineage, but it is done using the function mix (defined in (8e) because the left input lineage θ_x is not necessarily compatible with the right input lineage θ_y. Given this, Rule (7f) generates a coGroup on the join key first, and then, for each join key k, it groups the left and right join matches by θ_x and θ_y respectively, so that the output contains unique lineage key combinations, $(k, (\theta_x, \theta_y))$. Finally, Rule (7g) annotates each value of the input stream S_i with the empty lineage ().

We will prove next that the transformed query $\mathcal{T}_q[\![q]\!]$ is a homomorphism. We first prove that the generated sMap1 and sMap2 in $\mathcal{T}_q[\![q]\!]$ are homomorphisms, provided that their functional arguments are homomorphisms. More specifically, we prove the following judgments:

Lemma 1 (Transformed Term Judgments)**.**

$$\frac{\rho \vdash \text{sMap1}(f, X) : \Uparrow_{\otimes}}{\rho \vdash X : \Uparrow_{\oplus}, \ \rho[k : \square, v : \oplus] \vdash f(k, v) : \Uparrow_{\otimes}} \quad (9a)$$

$$\frac{\rho \vdash \text{sMap2}(f, X) : \uplus}{\rho \vdash X : \Uparrow_{\oplus}, \ \rho[k : \square, v : \oplus] \vdash f(k, v) : \uplus} \quad (9b)$$

$$\frac{\rho \vdash \text{sMap3}(f, X) : \uplus}{\rho \vdash X : \uplus} \quad (9c)$$

$$\frac{\rho \vdash \text{swap}(X) : \uplus}{\rho \vdash X : \uplus} \quad (9d)$$

$$\frac{\rho \vdash \text{mix}(X) : \Uparrow_{\uplus \times \uplus}}{\rho \vdash X : \Uparrow_{\uplus \times \uplus}} \quad (9e)$$

The proof of this Lemma is given in Appendix A.

Definition 3 (Query Merger Monoid)**.** *The query merger monoid $\mathcal{M}[\![q]\!]$ of a normalized query term q is defined as follows:*

$$\begin{aligned} \mathcal{M}[\![\text{reduce}(\oplus, \text{cMap}(f, e))]\!] &= \Uparrow_{\oplus} \\ \mathcal{M}[\![\text{cMap}(f, S_i)]\!] &= \uplus \\ \mathcal{M}[\![\text{cMap}(f, e)]\!] &= \Uparrow_{\otimes} \qquad \text{if } e \neq S_i \end{aligned}$$

where the monoid $\Uparrow_\otimes$ in the last equation comes from $\text{sMap1}(f,\ X) : \Uparrow_\otimes$ *in Equation (9a).*

Based on Lemma 1, we can now prove that the transformed query is a homomorphism:

Theorem 1 (Homomorphism)**.** *Any transformed term $\mathcal{T}_q[\![q]\!]$, where q is defined in Definition 2, is a homomorphism over the input streams, provided that each generated sMap1 and sMap2 term in $\mathcal{T}_q[\![q]\!]$ satisfies the premises in Judgment (9a) and (9b):*

$$\rho[S_1 : \uplus, \ldots, S_n : \uplus] \vdash \mathcal{T}_q[\![q]\!] :\ \mathcal{M}[\![q]\!] \tag{10}$$

The proof is given in Appendix A.

Definition 4 (Query Answer)**.** *The query answer $\mathcal{A}[\![q]\!]_x$ of the query q, defined in Definition 2, is a function over the current state $x = \mathcal{T}_q[\![q]\!]$ and is derived as follows:*

$$\begin{array}{l}
\mathcal{A}[\![\text{cMap}(f,\ S_i)]\!]_x = \pi_2(x) \\
\mathcal{A}[\![\text{reduce}(\oplus,\ \text{cMap}(f,\ e))]\!]_x = \text{reduce}(\oplus,\ \pi_2(x)) \\
\mathcal{A}[\![\text{cMap}(f,\ e)]\!]_x = \pi_2(\text{reduce}(\Uparrow_\otimes, T(x))) \qquad \textit{if } e \neq S_i \\
\mathcal{A}[\![q]\!]_x = \pi_2(\text{elem}(\text{reduce}(\Uparrow_\oplus, x))) \qquad \textit{otherwise}
\end{array}$$

where the monoid $\oplus$ in the last equation is $\mathcal{M}[\![q]\!]$, the monoid $\otimes$ in the second equation comes from $\text{sMap1}(f,\ X) : \Uparrow_\otimes$ *in Equation (9a), and*

$$\begin{array}{rcl}
\text{elem}(X) & = & \begin{cases} v & \textit{if } X = \{v\} \\ \text{error} & \textit{otherwise} \end{cases} \\
T(X) & = & \{\, (k, a) \,|\, ((k, \theta), a) \in X \,\}
\end{array}$$

The following theorem proves that $\mathcal{A}[\![q]\!]_x$ returns the same answer as q:

Theorem 2 (Correctness)**.** *The $\mathcal{A}[\![q]\!]_x$ over the state $x = \mathcal{T}_q[\![q]\!]$ returns the same result as the original query q, where q is defined in Definition 2:*

$$\mathcal{A}[\![q]\!]_x = q \qquad \textit{for } x = \mathcal{T}_q[\![q]\!] \tag{11}$$

The proof of this theorem is given in Appendix A.

8.1 Restricting Joins

Theorem 1 indicates that a query transformed by Algorithm 1 is a homomorphism if the cMap functional arguments in the query satisfy the premises in Judgment (9a) and (9b). But, consider the following join:

$$\begin{array}{l}
\text{cMap}(\lambda(k, (xs, ys)).\, \text{cMap}(\lambda x.\, \text{cMap}(\lambda y.\, \{(x, y)\},\ xs),\ ys), \\
\qquad \text{coGroup}(X, Y))
\end{array}$$

Our monoid inference system can not prove that the functional argument of this cMap is a homomorphism on both xs and ys. This is expected because this join could be a many-to-many join, which we know can not be a homomorphism. Since we have given up on handling many-to-many joins, we want to extend the monoid inference algorithm to always assume that every join is a one-to-one or one-to-many join and draw inferences based on this assumption.

We have already seen in Section 7 that if a term is invariant under change, such as an invariant data source, it is associated with the special monoid $\Box$. We can also use the annotation $\Box$ to denote certain functional dependencies, such as $\kappa \to \alpha$ on a bag X of type $\{(\kappa, \alpha)\}$, which indicates that the second component of a pair in X depends on the first. This dependency is captured by annotating $\text{groupBy}(X)$ with the monoid $\Uparrow_\Box$, which indicates that each group remains invariant under change, implying that the group-by key is also a unique key of X. That is, if $(k, s_1) \in X_1$ and $(k, s_2) \in X_2$, then $\{\, (k, s_1 \Box s_2) \,|\, (k, s_1) \in X_1,\ (k, s_2) \in X_2 \,\}$ in Definition (2) must have $s_1 = s_2$ so that $s_1 \Box s_2 = s_1$, otherwise it will be an error. This means that we can not have two different groups associated with the same key k. Given that a groupBy over a singleton gives a singleton group, each group returned from a groupBy is a singleton that remains invariant. A similar functional dependency can also apply to a join between X of type $\{(\kappa, \alpha)\}$ and Y of type $\{(\kappa, \beta)\}$, which is a $\text{coGroup}(X, Y)$ of type $\{(\kappa, (\{\alpha\}, \{\beta\}))\}$. To indicate that this join is one-to-one or one-to-many, we annotate $\text{coGroup}(X, Y)$ with the monoid $\Uparrow_{\Box \times \uplus}$, which enforces the constraint that the bag $\{\alpha\}$ in $\{(\kappa, (\{\alpha\}, \{\beta\}))\}$ be either empty or singleton, that is, at most one X element can be joined with Y over a key κ. This is because the X and Y values that are joined over the same key are merged with $\Box$ and $\uplus$, respectively, as indicated by $\Uparrow_{\Box \times \uplus}$. Therefore, to incorporate the assumption that all joins are one-to-one or one-to-many in the monoid inference system, we have to replace Judgment (5g) with the following judgment:

$$\frac{\rho \vdash \text{coGroup}(X, Y) : \Uparrow_{\Box \times \uplus}}{\rho \vdash X : \uplus,\ \rho \vdash Y : \uplus} \tag{12}$$

Given this judgment, the cMap input of our join example will be annotated with $\Uparrow_{\Box \times \uplus}$, which means that xs and ys will be annotated with $\Box$ and $\uplus$, respectively (ie, xs is invariant). Based on these annotations, the functional argument $\text{cMap}(\lambda x.\, \text{cMap}(\lambda y.\, \{(x, y)\},\ xs),\ ys)$ satisfies the premise in Judgment (9b) since it is annotated with $\uplus$.

8.2 Handling Iterative Queries

The approximate algorithm that processes iterations incrementally, presented in Section 2, requires an additional operation, called the diffusion operator $\widehat{\otimes}$, that satisfies the property:

$$T \otimes (T \mathbin{\widehat{\otimes}} \Delta T) = T \otimes (T \otimes \Delta T)$$

Our goal is to define $\widehat{\otimes}$ in terms of $\otimes$ in such a way that $T \mathbin{\widehat{\otimes}} \Delta T$ contains only those parts of $T \otimes \Delta T$ that changed from T.

Definition 5 (Diffusion)**.** *The diffusion $\widehat{\otimes}$ of a monoid $\otimes$ is defined as follows:*

$$\begin{array}{rcl}
\widehat{\Uparrow_\oplus} & = & \Downarrow_{\widehat{\oplus}} \qquad \widehat{\otimes \times \oplus} \qquad = \qquad \widehat{\otimes} \times \widehat{\oplus} \\
\widehat{\oplus} & = & \oplus \qquad \textit{for any other monoid}
\end{array}$$

where $\Downarrow_\otimes$ is a right-outer join defined as follows:

$$\begin{array}{rcl}
X \Downarrow_\otimes Y & \triangleq & \{\, (k, a \otimes b) \,|\, (k, a) \in X,\ (k, b) \in Y \,\} \\
& & \uplus\, \{\, (k, b) \,|\, (k, b) \in Y,\ k \notin \pi_1(X) \,\}
\end{array}$$

This $\widehat{\otimes}$ satisfies the desired property, $T \otimes (T \mathbin{\widehat{\otimes}} \Delta T) = T \otimes (T \otimes \Delta T)$ (proven in Appendix A).

9. NON-HOMOMORPHIC TERMS

Theorem 2 indicates that the terms generated by the transformations (7a) through (7h) are homomorphisms as long as the premises of the judgements in Lemma 1 are true. These premises indicate that the cMap functional arguments must be homomorphisms too. We want the operations that cannot be annotated with a monoid to be transformed so that

the non-homomorphic parts of the operation are pulled outwards from the query using rewrite rules. We achieve this with the help of kMap:

$$\text{kMap}(f, X) \triangleq \text{cMap}(\lambda(\theta, v).\, \{(\theta, f(v))\},\, X)$$

which is a cMap that propagates the lineage θ as is. In our framework, all non-homomorphic parts take the form of a kMap and are accumulated into one kMap using

$$\text{kMap}(f, \text{kMap}(g, X)) \;\rightarrow\; \text{kMap}(\lambda x.\, f(g(x)),\, X)$$

More specifically, we first split each non-homomorphic cMap to a composition of a kMap and a cMap so that the latter cMap is a homomorphism, and then we pull and merge kMaps. Consider the term $\text{cMap}(\lambda(\theta, v).\, \{(\theta', e)\},\, X)$, which creates a new lineage θ' from the old θ. In our framework, we find the largest subterms in the algebraic term e, namely $e_1, \ldots, e_n$, that are homomorphisms. This is accomplished by traversing the tree that represents the term e, starting from the root, and by checking if the node can be inferred to be a homomorphism. If it is, the node is replaced with a new variable. Thus, e is mapped to a term $f(e_1, \ldots, e_n)$, for some term f, and the terms $e_1, \ldots, e_n$ are replaced with variables when f is pulled outwards:

$$\begin{aligned}
&\text{cMap}(\lambda(\theta, v).\, \{(\theta', e)\},\, X)\\
&\quad\rightarrow \text{cMap}(\lambda(\theta, v).\, \{(\theta', f(e_1, \ldots, e_n))\},\, X)\\
&\quad\rightarrow \text{kMap}(\lambda(x_1, \ldots, x_n).\, f(x_1, \ldots, x_n),\\
&\qquad\qquad \text{cMap}(\lambda(\theta, v).\, \{(\theta', (e_1, \ldots, e_n))\},\, X))
\end{aligned}$$

The kMaps are combined and are pulled outwards from the query using the following rewrite rules:

$$\begin{aligned}
&\text{sMap2}(g, \text{kMap}(f, X)) \rightarrow \text{sMap2}(\lambda(k, v).\, g(k, f(v)),\, X)\\
&\text{sMap1}(g, \text{kMap}(f, X)) \rightarrow \text{sMap1}(\lambda(k, v).\, g(k, f(v)),\, X)\\
&\text{groupBy}(\text{kMap}(f, X)) \rightarrow \text{kMap}(\lambda s.\, \text{map}(f, s),\, \text{groupBy}(X))\\
&\text{coGroup}(\text{kMap}(f_x,\, X),\, \text{kMap}(f_y, Y))\\
&\quad\rightarrow \text{kMap}(\lambda(k, (xs, ys)).\, (k, (\text{map}(f_x, xs), \text{map}(f_y, ys))),\\
&\qquad\qquad \text{coGroup}(X, Y))
\end{aligned}$$

where $\text{map}(f, X) = \text{cMap}(\lambda x.\, \{f(x)\},\, X)$. Rewrite rules as these, when applied repeatedly, can pull out and combine the non-homomorphic parts of a query, leaving a homomorphism whose merge function can be derived from our annotation rules.

10. AN EXAMPLE

Consider again the algebraic term $Q(S_1, S_2)$, presented at the end of Section 5. If we apply the transformations in Equations (7a) through (7h), we derive the following term term $\mathcal{T}_q[\![Q(S_1, S_2)]\!]$, which propagates the join and group-by keys to the query output:

$$\begin{aligned}
&\text{sMap1}(\lambda(k, s).\, \{(k, avg(s))\},\\
&\quad \text{groupBy}(\text{swap}(\text{sMap2}(\lambda(j, (xs, ys)).\, g(xs, ys),\\
&\qquad \text{mix}(\text{coGroup}(\text{sMap3}(\lambda x.\, \{(x.B, x)\},\, S_1),\\
&\qquad\qquad \text{sMap3}(\lambda y.\, \{(y.C, y)\},\, S_2)))))))
\end{aligned} \tag{13}$$

where $avg(s) = \text{reduce}(+, s)/\text{reduce}(+, \text{cMap}(\lambda v.\, \{1\}, s))$ and $g(xs, ys) = \text{cMap}(\lambda x.\, \text{cMap}(\lambda y.\, \{(x.A, y.D)\},\, ys),\, xs)$. If we expand sMap1, sMap2, swap, and mix, then the transformed query is:

$$\begin{aligned}
&\{\, (\theta, (k, avg(s)))\\
&\;|\, (\theta, (k, s)) \in \text{groupBy}(\{\, ((k, (j, ((), ()))), v)\\
&\qquad\qquad\qquad\qquad |\, (j, (s_1, s_2)) \in join,\\
&\qquad\qquad\qquad\qquad\;\; ((), xs) \in \text{groupBy}(s_1),\\
&\qquad\qquad\qquad\qquad\;\; ((), ys) \in \text{groupBy}(s_2),\\
&\qquad\qquad\qquad\qquad\;\; (k, v) \in g(xs, ys)\,\})\,\}
\end{aligned}$$

where $join$ is

$$\text{coGroup}(\{\, (x.B, ((), x)) \,|\, x \in S_1\,\},\, \{\, (y.C, ((), y)) \,|\, y \in S_2\,\})$$

The output lineage θ is $(k, (j, ((), ())))$, where k and j are groupBy and coGroup keys. But, $\text{groupBy}(s_1) = \{((), \pi_2(s_1))\}$ and $\text{groupBy}(s_2) = \{((), \pi_2(s_2))\}$ since the key is $()$. Consequently, the transformed query becomes:

$$\begin{aligned}
&\{\, (\theta, (k, avg(s)))\\
&\;|\, (\theta, (k, s)) \in \text{groupBy}(\{\, ((k, (j, ((), ()))), v)\\
&\qquad\qquad\qquad\qquad |\, (j, (s_1, s_2)) \in join,\\
&\qquad\qquad\qquad\qquad\;\; (k, v) \in g(\pi_2(s_1), \pi_2(s_2))\,\})\,\}
\end{aligned}$$

We now check if the transformed query (13) is a homomorphism. Both coGroup inputs in (13) are sMap2 terms, annotated with $\uplus$, which means that, based on Equation (5g), the coGroup is annotated with $\Uparrow_{\square \times \uplus}$. From Equation (9b), the sMap2 operation is annotated with $\uplus$ as long as $g(xs, ys)$ is annotated with $\uplus$. Hence, the groupBy operation is annotated with $\Uparrow_{\uplus}$, based on Equation (5f). Unfortunately, based on Equation (8a), the sMap1 term is not a homomorphism as is, because $avg(s)$ is not a homomorphism over s. Based on the methods described in Section 9 that factor out non-homomorhic parts from terms, the sMap1 term is broken into two terms $a(h(S_1, S_2))$, where $h(S_1, S_2)$ is:

$$\begin{aligned}
&\text{sMap1}(\lambda(k, s).\, \{(k, (\text{reduce}(+, s),\\
&\qquad\qquad \text{reduce}(+, \text{cMap}(\lambda v.\, 1, s))))\}, \ldots)
\end{aligned}$$

This is equivalent to the homomorphism $h(S_1, S_2)$ given in Section 2.

In addition, from Definition 4, the answer function $a(x)$ is $\pi_2(\text{reduce}(\Uparrow_{\otimes}, T(x)))$. Therefore, the answer function $a(x)$, combined with the non-homomorphic part of the query is:

$$\begin{aligned}
&\text{cMap}(\lambda(k, (s, c)).\, \{(k, s/c)\},\\
&\qquad \text{reduce}(\Uparrow_{\square \times ((+) \times (+))},\, T(x)))
\end{aligned}$$

This is equivalent to the answer query given in Section 2. Finally, $h(S_1, S_2)$ is a homomorphism annotated with $\Uparrow_{\square \times ((+) \times (+))}$, because the reduce terms are annotated with $(+)$ (from Equation (5c)) and we have a pair of homomorphisms. This is equivalent to the merge function $X \otimes Y$ given in Section 2, implemented as a partitioned join combined with aggregation.

11. HANDLING DELETIONS

Our framework can be easily extended to handle deletion of existing data from an input stream, in addition to insertion of new data. To handle both insertions and deletions, a data stream S_i in our framework consists of an initial data set, followed by a continuous stream of incremental batches ΔS_i and a continuous stream of decremental batches $\Delta S_i'$ that arrive at regular time intervals Δt. Then, the stream data at time $t + \Delta t$ is $S_i \uplus \Delta S_i - \Delta S_i'$, where $X - Y$ is bag difference, which satisfies $(X \uplus Y) - Y = X$. (An element appears in $X - Y$ as many times as it appears in X, minus

the number of times it appears in Y.) This dual stream of updates, can be implemented as a single stream that contains updates tagged with + or -, to indicate insertion or deletion. Alternatively, a stream data source may monitor two separate directories, one for insertions and another for deletions, so that if a new file is created, it will be treated as a new batch of insertions or deletions, depending on the directory. In our framework, we require that $\Delta S'_i \subseteq S_i$, that is, we can only delete values that have already appeared in the stream in larger or equal multiplicities. This restriction implies that there is a bag W_i such that $S_i = W_i \uplus \Delta S'_i$. Without this restriction, it would be hard to diminish the query results on S_i by the results on ΔS_i to calculate the new results.

Definition 6 (Diminisher)**.** *The diminisher $\widetilde{\otimes}$ of a monoid $\otimes$ is defined as follows:*

$$\begin{array}{lcllcllcl} \widetilde{\uplus} & = & - & & \widetilde{+} & = & - & \widetilde{\square} & = & \square \\ \widetilde{\otimes \times \oplus} & = & \widetilde{\otimes} \times \widetilde{\oplus} & & \widetilde{\Uparrow_\oplus} & = & \Downarrow_{\widetilde{\oplus}} & & & \end{array}$$

where $\Downarrow_\otimes$ is a left-outer join defined as follows:

$$\begin{array}{rcl} X \Downarrow_\otimes Y & \triangleq & \{\,(k, a \otimes b) \,|\, (k,a) \in X,\ (k,b) \in Y,\ a \neq b\,\} \\ & & \uplus\, \{\,(k,a) \,|\, (k,a) \in X,\ k \notin \pi_1(Y)\,\} \end{array}$$

Note that a diminisher is not a monoid.

Theorem 3. *For all X, Y : $(X \otimes Y) \mathbin{\widetilde{\otimes}} Y = X$.*

The proof is given in Appendix A. For $X = \otimes_z$, it implies that $Y \mathbin{\widetilde{\otimes}} Y = \otimes_z$.

Theorem 4. *If $[S_1 : \uplus, \ldots, S_n : \uplus] \vdash f(\overline{S}) : \otimes$ and for all i : $\Delta S'_i \subseteq S_i$, then $f(\overline{S - \Delta S'}) = f(\overline{S}) \mathbin{\widetilde{\otimes}} f(\overline{\Delta S'})$.*

Proof. Since $\Delta S'_i \subseteq S_i$, then there must be W_i such that $S_i = W_i \uplus \Delta S'_i$. Thus, $S_i - \Delta S'_i = W_i \uplus \Delta S'_i - \Delta S'_i \Rightarrow W_i = S_i - \Delta S'_i$. Then,

$$\begin{array}{l} f(\overline{S}) \mathbin{\widetilde{\otimes}} f(\overline{\Delta S'}) = f(\overline{W \uplus \Delta S'}) \mathbin{\widetilde{\otimes}} f(\overline{\Delta S'}) \\ = (f(\overline{W}) \otimes f(\overline{\Delta S'})) \mathbin{\widetilde{\otimes}} f(\overline{\Delta S'}) = f(\overline{W}) = f(\overline{S - \Delta S'}) \end{array} \quad \square$$

This theorem indicates that, to process the deletions $\overline{\Delta S'}$, we can simply use the same methods as for insertions, but using the diminishing function $\widetilde{\otimes}$ for merging the state, instead of $\otimes$. That is, we can use the same functions h and a, but now the state is diminished as follows:

$$\begin{array}{l} \text{state} \leftarrow \text{state} \mathbin{\widetilde{\otimes}} h(\overline{\Delta S'}) \\ \textbf{return } a(\text{state}) \end{array}$$

For example, the join-groupby query used in Section 2 must now use the merging, $X \Downarrow_{(-)\times(-)} Y$, equal to:

```
select (θ, (sx−sy,cx−cy))
  from (θ,(sx,cx)) in X,
       (θ,(sy,cy)) in Y
union select x from x in X where π1(x) not in π1(Y)
```

12. IMPLEMENTATION

We have implemented our incremental processing framework using Apache MRQL [33] on top of Apache Spark Streaming [46]. The Spark streaming engine monitors the file directories used as stream sources in an MRQL query, and when a new file is inserted in one of these directories or the modification time of a file changes, it triggers the MRQL query processor to process the new files, based on the state derived from the previous step, and creates a new state.

We have introduced a new physical operator, called Incr, which is a stateful operator that updates a state. More specifically, every instance of this operation is annotated with a state number i and is associated with a state, state_i, of type α_i. The operation $\text{Incr}(i, s_0, F)$, where F is a state transition function of type $\alpha_i \to \alpha_i$ and s_0 is the initial state of type α_i, has the following semantics:

$$\begin{array}{l} \text{state}_i \leftarrow F(\text{state}_i) \\ \textbf{return } \text{state}_i \end{array}$$

with $\text{state}_i = s_0$, initially. Note that the states are preserved across the Inc calls and are modified by these calls.

Recall that, in our framework, we break a query q into a homomorphism h and an answer function a, and we evaluate the query incrementally using:

$$\begin{array}{l} \text{state}_1 \leftarrow \text{state}_1 \otimes h(\overline{\Delta S}) \\ \textbf{return } a(\text{state}_1) \end{array}$$

at every time interval, where, initially, $\text{state}_1 = \otimes_z$. This is implemented using the following physical plan:

$$a(\text{Incr}(1, \otimes_z, \lambda T.\ T \otimes h(\overline{\Delta S})))$$

For an iteration $\text{repeat}(\lambda X.\ f(X, \overline{\Delta S}), n, g(\overline{S}))$, we use the approximate solution, described in Section 2: we split f into a homomorphism h, for some monoid $\otimes$, and a function a, and we derive a diffusion operator $\widehat{\otimes}$:

$$\begin{array}{l} a(\text{Incr}(1, \otimes_z, \\ \quad \lambda T.\ T \otimes \text{repeat}(\lambda \Delta T.\ T \mathbin{\widehat{\otimes}} h(a(\Delta T), \overline{\Delta S}), \\ \qquad\qquad n-1, \\ \qquad\qquad T \mathbin{\widehat{\otimes}} h(g(\overline{\Delta S}), \overline{\Delta S})))) \end{array}$$

13. PERFORMANCE EVALUATION

The system described in this paper is available as part of the latest official MRQL release (MRQL 0.9.6). We have experimentally validated the effectiveness of our methods using four queries: groupBy, join-groupBy, k-means clustering, and PageRank. The platform used for our evaluations is a small cluster of 9 Linux servers, connected through a Gigabit Ethernet switch. Each server has 4 Xeon cores at 3.2GHz with 4GB memory. For our experiments, we used Hadoop 2.2.0 (Yarn) and Spark 1.6.0. The cluster frontend was used exclusively as a NameNode/ResourceManager, while the rest 8 compute nodes were used as DataNodes/NodeManagers. For our experiments, we used all the available 32 cores of the compute nodes for Spark tasks.

The data streams used by the first 3 queries (groupBy, join-groupBy, and k-means clustering) consist of a large set of initial data, which is used to initialize the state, followed by a sequence of 9 equal-size batches of data (the increments). The experiments were repeated for increments of size 10K, 20K, 30K, and 40K tuples, always starting with a fresh state (constructed from the initial data only). The performance results are shown in Figure 3. The x-axis represents the time points Δt when we get new batches of data in the stream. At time $0\Delta t$, we have the processing of the initial data and the construction of the initial state. Then, the 9 increments arrive at the time points $1\Delta t$ through $9\Delta t$. The y-axis is the query execution time, and there are 4 plots, one for each increment size.

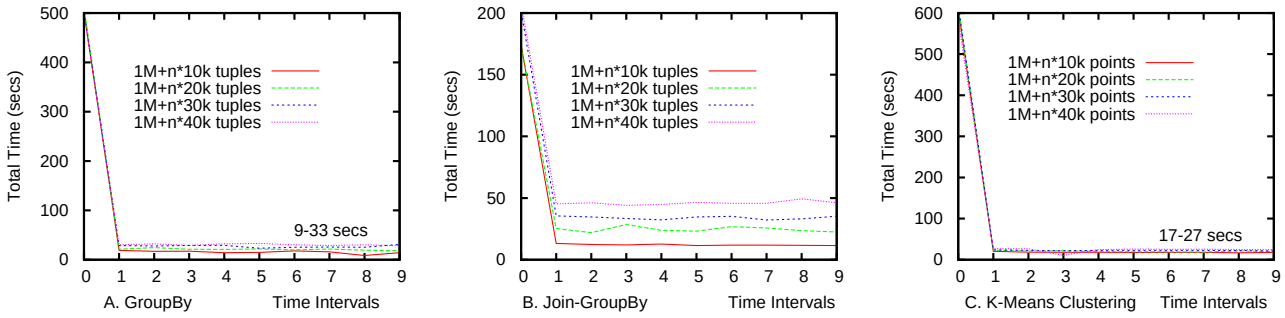

Figure 3: Incremental Query Evaluation of GroupBy, Join-GroupBy, and K-Means Clustering

The join-groupBy and the k-means clustering queries are given in Section 2. The groupBy query is 'select (x,avg(y)) from (x,y) in S group by x'. The datasets used for both the groupBy and join-groupBy queries consist of pairs of random integers between 0 and 10000. The groupBy initial dataset has size 1M tuples, while the two join-groupBy inputs have size 100K tuples. The datasets used for the k-means query consist of random (X, Y) points in 4 squares that have X in $[2, 4]$ or $[6, 8]$ and Y in $[2, 4]$ or $[6, 8]$. Thus, the 4 centroids are expected to be $(3, 3)$, $(3, 7)$, $(7, 3)$, and $(7, 7)$. This also means that the state contains 4 centroids only. The initial dataset for k-means contains 1M points and the k-means query uses 10 iteration steps. The k-means incremental program uses the approximate solution described in Section 2 based on approximate floating point equality.

From Figure 3, we can see that processing incremental batches of data can give an order of magnitude speed-up compared to processing all the data each time. Furthermore, the time to process each increment does not substantially increase through time, despite that the state grows with new data each time (in the case of the groupBy and join-groupBy queries). This happens because merging states is done with a partitioned join (implemented as a coGroup in Spark) so that the new state created by coGroup is already partitioned by the join key and is ready to be used for the next coGroup to handle the next increment. Consequently, only the results of processing the new data, which are typically smaller than the state, are shuffled across the worker nodes before coGroup. This makes the incremental processing time largely independent of the state size in most cases since data shuffling is the most prevalent factor in main-memory distributed processing systems.

We have also evaluated our system on a PageRank query on random graphs generated by the R-MAT algorithm [9] using the Kronecker graph generator parameters: a=0.30, b=0.25, c=0.25, and d=0.20. The PageRank query expressed in MRQL is a simple self-join on the graph, which is optimized into a single group-by operation. (The MRQL PageRank query is given in [18].) The PageRanks were calculated in 10 iteration steps. This time, the initial graph used in our evaluations had size 5K nodes with 50K edges and the four different increments had sizes 100 nodes with 100, 200, 300, and 400 edges, respectively. The performance results are shown in Figure 4. We believe that the reason that we did not get a flat line for our incremental PageRank implementation was due to the lack of sufficient memory tuning. In Spark, cached RDDs and D-Streams are not automatically garbage-collected; instead, Spark leaves it to

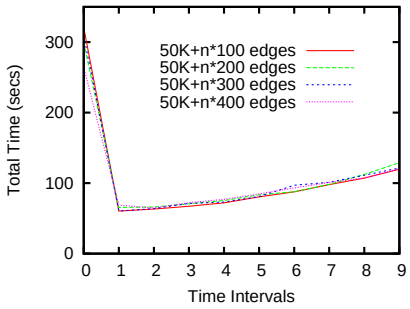

Figure 4: Incremental Query Evaluation of PageRank

the programmer to cache the queued operations in memory if their results are used multiple times, and to dispose the cache later. This is reminiscent to legacy programming languages, such as C, that did not support GC. Removing all memory leaks is very crucial to continuous stream processing, and, if not done properly, it will eventually cause the system to run out of memory. We are planning to fine-tune our streaming engine to remove all these memory leaks.

14. CONCLUSION AND FUTURE WORK

We have presented a general framework for translating batch MRQL queries to incremental DSPE programs. In contrast to other systems, our methods are completely automated and are formally proven to be correct. In addition to incremental query processing on a DSPE platform, our framework can also be used on a batch distributed system to process data larger than the available distributed memory, by processing these data incrementally, in batches that can fit in memory. Although our methods are described using the unconventional MRQL algebra, instead of the nested relational algebra, we believe that many other similar query systems can use our framework by simply translating their algebraic operators to the MRQL operators, then, using our framework as is, and, finally, translating the resulting operations back to their own algebra.

As a future work, we are planning to use an in-memory key-value store to update the state in place, instead of replacing the entire state with a new one every time we process a new batch of data. Spark Streaming, in fact, already supports an operation updateStateByKey operation, which

allows to maintain an arbitrary state while continuously updating it with new data. We believe that this will considerably improve the performance of queries that require large states, such as PageRank. Finally, in the near future, we are planning to add support for more DSPE platforms in Incremental MRQL, such as Storm and Flink Streaming.

Acknowledgments: This work is supported in part by the National Science Foundation under the grant CCF-1117369. Some of our performance evaluations were performed at the Chameleon cloud computing infrastructure, supported by NSF, `https://www.chameleoncloud.org/`.

15. REFERENCES

[1] D. J. Abadi, D. Carney, U. Cetintemel, et al. Aurora: A New Model and Architecture for Data Stream Management. In *VLDB Journal*, 12(2):120–139, 2003.

[2] U. A. Acar, G. E. Blelloch, M. Blume, R. Harper, and K. Tangwongsan. An Experimental Analysis of Self-Adjusting Computation. In *ACM Trans. Prog. Lang. Sys., 32(1):3:153, 2009.*

[3] U. A. Acar and Y. Chen. Streaming Big Data with Self-Adjusting Computation. In *Workshop on Data Driven Functional Programming (DDFP)*, 2013.

[4] B. Babcock, S. Babu, M. Datar, R. Motwani, and J. Widom. Models and Issues in Data Stream Systems. In *Symposium on Principles of Database Systems (PODS)*, pages 1–16, 2002.

[5] O. Benjelloun, A. D. Sarma, A. Halevy, and J. Widom. ULDBs: Databases with Uncertainty and Lineage. In *International Conference on Very Large Data Bases (VLDB)*, pages 953–964, 2006.

[6] D. Bhagwat, L. Chiticariu, W. C. Tan, and G. Vijayvargiya. An Annotation Management System for Relational Databases. In *International Conference on Very Large Data Bases (VLDB)*, pages 900–911, 2004.

[7] P. Bhatotia, A. Wieder, R. Rodrigues, U. A. Acar, and R. Pasquin. Incoop: Mapreduce for Incremental Computations. In *ACM Symposium on Cloud Computing (SoCC)*, 2011.

[8] Y. Cai, P. G. Giarrusso, T. Rendel, and K. Ostermann. A Theory of Changes for Higher-Order Languages. Incrementalizing λ-Calculi by Static Differentiation. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 145-155, 2014.

[9] D. Chakrabarti, Y. Zhan, and C. Faloutsos. R-MAT: A Recursive Model for Graph Mining. In *Fourth SIAM International Conference on Data Mining (SDM)*, pages 442–446, 2004.

[10] B. Chandramouli, J. Goldstein, M. Barnett, R. DeLine, D. Fisher, J. C. Platt, J. F. Terwilliger, J. Wernsing. Trill: A High-Performance Incremental Query Processor for Diverse Analytics. In *International Conference on Very Large Data Bases (VLDB)*, pages 401–412, 2014.

[11] S. Chandrasekaran, O. Cooper, A. Deshpande, M. J. Franklin, J. M. Hellerstein, W. Hong, S. Krishnamurthy, S. Madden, V. Raman, F. Reiss, and M. Shah. TelegraphCQ: Continuous Data flow Processing for an Uncertain World. In *Conference on Innovative Data System Research (CIDR)*, 2003.

[12] T. Condie, N. Conway, P. Alvaro, J. M. Hellerstein, K. Elmeleegy, and R. Sears. Mapreduce Online. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 10(4), 2010.

[13] Y. Cui and J. Widom. Lineage Tracing for General Data Warehouse Transformations. In *International Conference on Very Large Data Bases (VLDB)*, pages 471–480, 2001.

[14] J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. In *Symposium on Operating System Design and Implementation (OSDI)*, 2004.

[15] P. Desikan, N. Pathak, J. Srivastava, and V. Kumar Incremental PageRank Computation on Evolving Graphs. In *International conference on World Wide Web (WWW)*, pages 1094–1095, 2005.

[16] L. Fegaras. Supporting Bulk Synchronous Parallelism in Map-Reduce Queries. In *International Workshop on Data Intensive Computing in the Clouds (DataCloud)*, 2012.

[17] L. Fegaras, C. Li, U. Gupta, and J. J. Philip. XML Query Optimization in Map-Reduce. In *International Workshop on the Web and Databases (WebDB)*, 2011.

[18] L. Fegaras, C. Li, and U. Gupta. An Optimization Framework for Map-Reduce Queries. In *International Conference on Extending Database Technology (EDBT)*, pages 26–37, 2012.

[19] L. Fegaras and D. Maier. Optimizing Object Queries Using an Effective Calculus. In *ACM Transactions on Database Systems (TODS)*, 25(4):457–516, 2000.

[20] Apache Flink. `http://flink.apache.org/`.

[21] Apache Giraph. `http://giraph.apache.org/`.

[22] A. Ghoting, R. Krishnamurthy, E. Pednault, B. Reinwald, V. Sindhwani, S. Tatikonda, Y. Tian, and S. Vaithyanathan. SystemML: Declarative Machine Learning on MapReduce. In *IEEE International Conference on Data Engineering (ICDE)*, pages 231–242, 2011.

[23] A. Gupta and I. S. Mumick. Maintenance of Materialized Views: Problems, Techniques, and Applications. In *IEEE Bulletin on Data Engineering*, 18(2):145–157, 1995.

[24] Apache Hadoop. `http://hadoop.apache.org/`.

[25] Apache Hama. `http://hama.apache.org/`.

[26] Apache Hive. `http://hive.apache.org/`.

[27] D. Logothetis, C. Olston, B. Reed, K.C. Webb, and K. Yocum. Stateful Bulk Processing for Incremental Analytics. In *ACM Symposium on Cloud Computing (SoCC)*, 2010.

[28] Y. Low, J. Gonzalez, A. Kyrola, D. Bickson, C. Guestrin, and J. M. Hellerstein. Distributed GraphLab: A Framework for Machine Learning and Data Mining in the Cloud. In *Proceedings of the VLDB Endowment*, 5(8):716-727, 2012.

[29] G. Malewicz, M. H. Austern, A. J.C Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski. Pregel: a System for Large-Scale Graph Processing. In *ACM symposium on Principles of Distributed Computing (PODC)*, 2009.

[30] F. McSherry, D. G. Murray, R. Isaacs, and M. Isard. Differential Dataflow. In *Conference on Innovative*

Data System Research (CIDR), 2013.

[31] S. R. Mihaylov, Z. G. Ives, and S. Guha. REX: Recursive, Delta-Based Data-Centric Computation. In *Proceedings of the VLDB Endowment*, 5(11):1280-1291, 2012.

[32] D. G. Murray, F. McSherry, R. Isaacs, M. Isard, P. Barham, and M. Abadi. Naiad: a Timely Dataflow System. In *ACM Symposium on Operating Systems Principles (SOSP)*, 2013.

[33] Apache MRQL (incubating). `http://mrql.incubator.apache.org/`.

[34] C. Olston, B. Reed, U. Srivastava, R. Kumar, and A. Tomkins. Pig Latin: a not-so-Foreign Language for Data Processing. In *ACM SIGMOD International Conference on Management of Data*, pages 1099-1110, 2008.

[35] D. Peng and F. Dabek. Large-scale Incremental Processing using Distributed Transactions and Notifications. In *Symposium on Operating System Design and Implementation (OSDI)*, 2010.

[36] R. Power and J. Li. Piccolo: Building Fast, Distributed Programs with Partitioned Tables. In *Symposium on Operating System Design and Implementation (OSDI)*, 2010.

[37] Ramalingam, G. and Reps, T. A Categorized Bibliography on Incremental Computation. In *Principles of Programming Languages (POPL)*, 1993, pp. 502510.

[38] A. Shinnar, D. Cunningham, B. Herta, and V. Saraswat. M3R: Increased Performance for In-Memory Hadoop Jobs. In *Proceedings of the VLDB Endowment*, 5(12):1736-1747, 2012.

[39] Apache Samza. `http://samza.apache.org/`

[40] Apache Spark. `http://spark.apache.org/`.

[41] Apache Storm: A System for Processing Streaming Data in Real Time. `http://hortonworks.com/hadoop/storm/`.

[42] Apache S4 (incubating): A Distributed Stream Computing Platform. `http://incubator.apache.org/s4/`.

[43] K. Tangwongsan, M. Hirzel, S. Schneider, and K.-L. Wu. General Incremental Sliding-Window Aggregation. In *Proceedings of the VLDB Endowment*, 8(7):702-713, 2015.

[44] L. G. Valiant. A Bridging Model for Parallel Computation. In *Communications of the ACM (CACM)*, 33(8):103-111, August 1990.

[45] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica. Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2012.

[46] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker, and I. Stoica. Discretized Streams: Fault-Tolerant Streaming Computation at Scale. In *Symposium on Operating Systems Principles (SOSP)*, 2013.

[47] Y. Zhang, S. Chen, Q. Wang, and G. Yu. i2 MapReduce: Incremental MapReduce for Mining Evolving Big Data. In *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 27(7):1906-1919, 2015.

APPENDIX

A. PROOFS

Proof of Equation (3). We will use structural induction. The inductive step

$$\{\,(k,v)\,|\,(k,s)\in \text{groupBy}(X\uplus Y),\, v\in s\,\} \;=\; X\uplus Y$$

can be proven as follows. Let $GX=\text{groupBy}(X)$ and $GY=\text{groupBy}(Y)$. Then, for

$$\begin{aligned} M(X,Y) \;&=\{\,(k,a)\,|\,(k,a)\in X,\,(k,b)\in Y\,\} \\ &\quad\uplus\{\,(k,a)\,|\,(k,a)\in X,\,k\notin\pi_1(Y)\,\} \end{aligned} \tag{14}$$

we have, $M(GX,GY)=GX$, because for each group-by key k, there is at most one $(k,a)\in GX$ and at most one $(k,b)\in GY$. Then, we have:

$$\begin{aligned}
&\{\,(k,v)\,|\,(k,s)\in\text{groupBy}(X\uplus Y),\,v\in s\,\}\\
&=\{\,(k,v)\,|\,(k,s)\in(GX\Uparrow_{\uplus}GY),\,v\in s\,\}\\
&=\{\,(k,v)\,|\,(k,s)\in(\{\,(k,a\uplus b)\,|\,(k,a)\in GX,\,(k,b)\in GY\,\}\\
&\qquad\uplus\{\,(k,a)\,|\,(k,a)\in GX,\,k\notin\pi_1(GY)\,\}\\
&\qquad\uplus\{\,(k,b)\,|\,(k,b)\in GY,\,k\notin\pi_1(GX)\,\}),\,v\in s\,\}\\
&=\{\,(k,v)\,|\,(k,a)\in GX,\,(k,b)\in GY,\,v\in(a\uplus b)\,\}\\
&\quad\uplus\{\,(k,v)\,|\,(k,a)\in GX,\,k\notin\pi_1(GY),v\in a\,\}\\
&\quad\uplus\{\,(k,v)\,|\,(k,b)\in GY,\,k\notin\pi_1(GX),v\in b\,\}\\
&=\{\,(k,v)\,|\,(k,a)\in GX,\,(k,b)\in GY,\,v\in a\,\}\\
&\quad\uplus\{\,(k,v)\,|\,(k,a)\in GX,\,(k,b)\in GY,\,v\in b\,\}\\
&\quad\uplus\{\,(k,v)\,|\,(k,a)\in GX,\,k\notin\pi_1(GY),v\in a\,\}\\
&\quad\uplus\{\,(k,v)\,|\,(k,b)\in GY,\,k\notin\pi_1(GX),v\in b\,\}\\
&=\{\,(k,v)\,|\,(k,a)\in M(GX,GY),\,v\in a\,\}\\
&\quad\uplus\{\,(k,v)\,|\,(k,b)\in M(GY,GX),\,v\in b\,\}\\
&=\{\,(k,v)\,|\,(k,a)\in GX,\,v\in a\,\}\\
&\quad\uplus\{\,(k,v)\,|\,(k,b)\in GY,\,v\in b\,\}\\
&=X\uplus Y\qquad\textit{(from induction hypothesis)}
\end{aligned}$$

□

Proof of Lemma 1 (Transformed Term Judgments). Using Equations (8a) and (2), Judgement (9a) can be proven as follows:

$$\begin{aligned}
&\text{sMap1}(f,X\Uparrow_{\oplus}Y)\\
&=\{\,((\theta,k),b)\,|\,((\theta,k),a)\in(X\Uparrow_{\oplus}Y),\,b\in f(k,a)\,\}\\
&=\{\,((\theta,k),b)\,|\,((\theta,k),x)\in X,\,((\theta,k),y)\in Y,\,b\in f(k,x\oplus y)\,\}\uplus\cdots\\
&=\{\,((\theta,k),b)\,|\,((\theta,k),x)\in X,\,((\theta,k),y)\in Y,\\
&\qquad b\in(f(k,x)\Uparrow_{\otimes}f(k,y))\,\}\uplus\cdots\\
&=\{\,((\theta,k),n\otimes m)\,|\,((\theta,k),x)\in X,\,((\theta,k),y)\in Y,\\
&\qquad n\in f(k,x),\,m\in f(k,y)\,\}\uplus\cdots\\
&=\{\,((\theta,k),n\otimes m)\,|\,((\theta,k),x)\in X,\,n\in f(k,x),\\
&\qquad((\theta,k),y)\in Y,\,m\in f(k,y)\,\}\uplus\cdots\\
&=\{\,((\theta,k),n\otimes m)\,|\,((\theta,k),n)\in\text{sMap1}(f,X),\\
&\qquad((\theta,k),m)\in\text{sMap1}(f,Y)\,\}\uplus\cdots\\
&=\text{sMap1}(f,X)\Uparrow_{\otimes}\text{sMap1}(f,Y)
\end{aligned}$$

Using Equations (8b), (2), and (8d), Judgement (9b) can be

proven as follows:

$$
\begin{aligned}
&\text{sMap2}(f, X \Uparrow_{\oplus} Y)\\
&= \text{sMap2}(f, \{\, ((k,\theta), a \oplus b) \mid ((k,\theta), a) \in X,\ ((k,\theta), b) \in Y \,\}\\
&\quad \uplus \cdots)\\
&= \{\, (k', ((k,\theta), d)) \\
&\quad \mid ((k,\theta), c) \in (\{\, ((k,\theta), a \oplus b) \mid ((k,\theta), a) \in X,\\
&\qquad\qquad ((k,\theta), b) \in Y \,\} \uplus \cdots),\\
&\quad (k', d) \in f(k, c) \,\}\\
&= \{\, (k', ((k,\theta), d)) \mid ((k,\theta), a) \in X,\ ((k,\theta), b) \in Y,\\
&\qquad (k', d) \in f(k, a \oplus b) \,\} \uplus \cdots\\
&= \{\, (k', ((k,\theta), d)) \mid ((k,\theta), a) \in X,\ ((k,\theta), b) \in Y,\\
&\qquad (k', d) \in (f(k, a) \uplus f(k, b)) \,\} \uplus \cdots\\
&= \{\, (k', ((k,\theta), d)) \mid ((k,\theta), a) \in X,\ ((k,\theta), b) \in Y,\\
&\qquad (k', d) \in f(k, a) \,\}\\
&\quad \uplus \{\, (k', ((k,\theta), d)) \mid ((k,\theta), a) \in X,\ ((k,\theta), b) \in Y,\\
&\qquad (k', d) \in f(k, b) \,\} \uplus \cdots\\
&= (\{\, (k', ((k,\theta), d)) \mid ((k,\theta), a) \in X,\ ((k,\theta), b) \in Y,\\
&\qquad (k', d) \in f(k, a) \,\}\\
&\quad \uplus \{\, (k', ((k,\theta), d)) \mid ((k,\theta), a) \in X,\ (k,\theta) \notin \pi_1(Y),\\
&\qquad (k', d) \in f(k, a) \,\})\\
&\quad \uplus (\{\, (k', ((k,\theta), d)) \mid ((k,\theta), a) \in X,\ ((k,\theta), b) \in Y,\\
&\qquad (k', d) \in f(k, b) \,\}\\
&\qquad \uplus \{\, (k', ((k,\theta), d)) \mid ((k,\theta), b) \in Y,\ (k,\theta) \notin \pi_1(X),\\
&\qquad (k', d) \in f(k, b) \,\})\\
&= \{\, (k', ((k,\theta), d)) \mid ((k,\theta), a) \in X,\ (k', d) \in f(k, a) \,\})\\
&\quad \uplus \{\, (k', ((k,\theta), d)) \mid ((k,\theta), b) \in Y,\ (k', d) \in f(k, b) \,\}\\
&= \text{sMap2}(f, X) \uplus \text{sMap2}(f, Y)
\end{aligned}
$$

Judgement (9c) is a direct consequence of Judgement (5e). Judgement (9e) can be proven as follows:

$$
\begin{aligned}
&\text{mix}(X \Uparrow_{\uplus\times\uplus} Y)\\
&= \{\, ((k, (\theta_x, \theta_y)), (xs, ys)) \mid (k, (s_1, s_2)) \in (X \Uparrow_{\uplus\times\uplus} Y),\\
&\qquad (\theta_x, xs) \in \text{groupBy}(s_1),\ (\theta_y, ys) \in \text{groupBy}(s_2) \,\}\\
&= \{\, ((k, (\theta_x, \theta_y)), (xs, ys)) \mid (k, (s_1^1, s_2^1)) \in X,\ (k, (s_1^2, s_2^2)) \in Y,\\
&\qquad (\theta_x, xs) \in \text{groupBy}(s_1^1 \uplus s_1^2),\\
&\qquad (\theta_y, ys) \in \text{groupBy}(s_2^1 \uplus s_2^2) \,\} \uplus \cdots\\
&= \{\, ((k, (\theta_x, \theta_y)), (xs_1 \uplus xs_2, ys_1 \uplus ys_2))\\
&\quad \mid (k, (s_1^1, s_2^1)) \in X,\ (k, (s_1^2, s_2^2)) \in Y,\\
&\qquad (\theta_x, xs_1) \in \text{groupBy}(s_1^1),\ (\theta_x, xs_2) \in \text{groupBy}(s_1^2),\\
&\qquad (\theta_y, ys_1) \in \text{groupBy}(s_2^1),\ (\theta_y, ys_2) \in \text{groupBy}(s_2^2) \,\} \uplus \cdots\\
&= \{\, ((k, (\theta_x, \theta_y)), (xs_1 \uplus xs_2, ys_1 \uplus ys_2))\\
&\quad \mid ((k, (\theta_x, \theta_y)), (xs_1, ys_1)) \in \text{mix}(X),\\
&\qquad ((k, (\theta_x, \theta_y)), (xs_2, ys_2)) \in \text{mix}(Y) \,\} \uplus \cdots\\
&= \text{mix}(X) \Uparrow_{\uplus\times\uplus} \text{mix}(Y)
\end{aligned}
$$

□

Theorem 1 (Homomorphism)**.** *Any transformed term $\mathcal{T}_q[\![q]\!]$, where q is defined in Definition 2, is a homomorphism over the input streams, provided that each generated sMap1 and sMap2 term in $\mathcal{T}_q[\![q]\!]$ satisfies the premises in Judgment (9a) and (9b):*

$$\rho[S_1 : \uplus, \ldots, S_n : \uplus] \vdash \mathcal{T}_q[\![q]\!] : \mathcal{M}[\![q]\!] \tag{10}$$

Proof. Using induction, we can first prove that $\mathcal{T}_c[\![c]\!] : \uplus$, $\mathcal{T}_e[\![\text{groupBy}(c)]\!] : \Uparrow_{\uplus}$, and $\mathcal{T}_e[\![\text{coGroup}(c_1, c_2)]\!] : \Uparrow_{\uplus\times\uplus}$, for all c, c_1, c_2. We have $\mathcal{T}_c[\![\text{cMap}(f, S_i)]\!] = \text{sMap3}(f, S_i) : \uplus$ from Judgment (9c), and $\mathcal{T}_c[\![\text{cMap}(f, e)]\!] = \text{sMap2}(f, \mathcal{T}_e[\![e]\!]) : \uplus$ from Judgment (9b), because $\mathcal{T}_q[\![e]\!]$ is either $\Uparrow_{\uplus}$ or $\Uparrow_{\uplus\times\uplus}$. Similarly, $\mathcal{T}_e[\![\text{groupBy}(c)]\!] = \text{groupBy}(\text{swap}(\mathcal{T}_c[\![c]\!])) : \Uparrow_{\uplus}$ from Judgments (5f) and (9d), and $\mathcal{T}_e[\![\text{coGroup}(c_1, c_2)]\!]$ $= \text{mix}(\text{coGroup}(\mathcal{T}_c[\![c_1]\!], \mathcal{T}_c[\![c_2]\!])) : \Uparrow_{\uplus\times\uplus}$ from Judgments (5g) and (9e). Then, $\mathcal{T}_q[\![\text{cMap}(f, e)]\!] = \text{sMap1}(f, \mathcal{T}_e[\![e]\!]) : \Uparrow_{\otimes}$ from Judgment (9a), and $\mathcal{T}_q[\![\text{reduce}(\oplus, \text{cMap}(f, e))]\!] = \text{reduce}(\Uparrow_{\oplus}$ $, \text{sMap1}(f, \mathcal{T}_e[\![e]\!])) : \Uparrow_{\oplus}$ from Judgments (5c) and (9a). □

We will prove Theorem 2 using the following lemma:

Lemma 2. *The following equations apply to terms in Definition 2 for a function f that satisfies the premise of Equation (9b):*

$$G_1(\text{groupBy}(\text{swap}(X))) = \text{groupBy}(T'(X)) \tag{15a}$$

$$G_2(\text{mix}(\text{coGroup}(X, Y))) = \text{coGroup}(T'(X), T'(Y)) \tag{15b}$$

$$T'(\text{sMap2}(f, X)) = \text{cMap}(f, G_1(X)),\ \textit{if } X : \Uparrow_{\uplus} \tag{15c}$$

$$T'(\text{sMap2}(f, X)) = \text{cMap}(f, G_2(X)),\ \textit{if } X : \Uparrow_{\uplus\times\uplus} \tag{15d}$$

where T' removes the lineage θ, while G_1 and G_2, in addition to removing θ, reconstruct the groupBy and coGroup result:

$$T'(X) = \{\, (k, a) \mid (k, (\theta, a)) \in X \,\} \tag{16a}$$

$$G_1(X) = \text{groupBy}(\{\, (k, a) \mid ((k,\theta), s) \in X,\ a \in s \,\}) \tag{16b}$$

$$\begin{aligned} G_2(X) = \text{coGroup}(&\{\, (k, a) \mid ((k,\theta), (s_1, s_2)) \in X,\ a \in s_1 \,\},\\ &\{\, (k, b) \mid ((k,\theta), (s_1, s_2)) \in X,\ b \in s_2 \,\}) \end{aligned} \tag{16c}$$

Proof. Proof of Equation (15a) using Equation (3):

$$
\begin{aligned}
&G_1(\text{groupBy}(\text{swap}(X)))\\
&= \text{groupBy}(\{\, (k, a) \mid ((k,\theta), s) \in \text{groupBy}(\text{swap}(X)),\ a \in s \,\})\\
&= \text{groupBy}(T(\{\, ((k,\theta), a) \mid ((k,\theta), s) \in \text{groupBy}(\text{swap}(X)),\\
&\qquad\qquad a \in s \,\}))\\
&= \text{groupBy}(T(\text{swap}(X)))\\
&= \text{groupBy}(T'(X))
\end{aligned}
$$

Proof of Equation (15b):

$$
\begin{aligned}
&G_2(\text{mix}(\text{coGroup}(X, Y)))\\
&= \text{coGroup}(\{\, (k, a) \mid ((k,\theta), (s_1, s_2)) \in \text{mix}(\text{coGroup}(X, Y)),\\
&\qquad\qquad a \in s_1 \,\},\\
&\qquad \{\, (k, b) \mid ((k,\theta), (s_1, s_2)) \in \text{mix}(\text{coGroup}(X, Y)),\\
&\qquad\qquad b \in s_2 \,\})\\
&= \text{coGroup}(\{\, (k, a) \mid (k, (q_1, q_2)) \in \text{coGroup}(X, Y),\\
&\qquad\qquad (\theta_x, xs) \in \text{groupBy}(q_1),\\
&\qquad\qquad (\theta_y, ys) \in \text{groupBy}(q_2),\ a \in xs \,\},\\
&\qquad \{\, (k, b) \mid (k, (q_1, q_2))) \in \text{coGroup}(X, Y),\\
&\qquad\qquad (\theta_x, xs) \in \text{groupBy}(q_1),\\
&\qquad\qquad (\theta_y, ys) \in \text{groupBy}(q_2),\ b \in ys \,\})\\
&= \text{coGroup}(\{\, (k, a) \mid (k, (q_1, q_2)) \in \text{coGroup}(X, Y),\\
&\qquad\qquad a \in \pi_2(q_1) \,\},\\
&\qquad \{\, (k, b) \mid (k, (q_1, q_2)) \in \text{coGroup}(X, Y),\\
&\qquad\qquad b \in \pi_2(q_2) \,\})\\
&= \text{coGroup}(\{\, (k, a) \mid (k, (s_1, s_2)) \in \text{coGroup}(T'(X), T'(Y)),\\
&\qquad\qquad a \in s_1 \,\},\\
&\qquad \{\, (k, b) \mid (k, (s_1, s_2)) \in \text{coGroup}(T'(X), T'(Y)),\\
&\qquad\qquad b \in s_2 \,\})\\
&= \text{coGroup}(T'(X), T'(Y))
\end{aligned}
$$

It is easy to prove that, for an f that satisfies the premise of Equation (9b), we have $\text{cMap}(f, X) : \uplus$ for $X : \Uparrow_{\uplus}$. Then, we can prove Equation (15c) for $X \Uparrow_{\uplus} Y$ using Equation (9b)

and induction:

$$\begin{aligned}
&\text{cMap}(f, G_1(X \Uparrow_{\uplus} Y)) \\
&= \text{cMap}(f, (G_1(X) \Uparrow_{\uplus} G_1(Y))) \\
&= \text{cMap}(f, G_1(X)) \uplus \text{cMap}(f, G_1(Y)) \\
&= T'(\text{sMap2}(f, X)) \uplus T'(\text{sMap2}(f, Y)) \\
&= T'(\text{sMap2}(f, X \Uparrow_{\uplus} Y))
\end{aligned}$$

The proof of Equation (15d) is similar. □

Theorem 2 (Correctness)**.** *The $\mathcal{A}[\![q]\!]_x$ over the state $x = \mathcal{T}_q[\![q]\!]$ returns the same result as the original query q, where q is defined in Definition 2:*

$$\mathcal{A}[\![q]\!]_x = q \qquad \text{for } x = \mathcal{T}_q[\![q]\!] \tag{11}$$

Proof. To prove Equation (11), we first prove $T'(\mathcal{T}_c[\![c]\!]) = c$ from Equations (7g) and (7h). If $c = \text{cMap}(f, S_i)$, then

$$\begin{aligned}
&T'(\mathcal{T}_c[\![\text{cMap}(f, S_i)]\!]) \\
&= T'(\text{sMap3}(f, S_i)) \\
&= \text{cMap}(f, S_i)
\end{aligned}$$

If $e = \text{cMap}(f, \text{groupBy}(c))$, then using Equations (7h), (7e), (15a) and (15c):

$$\begin{aligned}
&T'(\mathcal{T}_c[\![\text{cMap}(f, \text{groupBy}(c))]\!]) \\
&= T'(\text{sMap2}(f, \mathcal{T}_e[\![\text{groupBy}(c)]\!])) \\
&= T'(\text{sMap2}(f, \text{groupBy}(\text{swap}(\mathcal{T}_c[\![c]\!])))) \\
&= \text{cMap}(f, G_1(\text{groupBy}(\text{swap}(\mathcal{T}_c[\![c]\!])))) \\
&= \text{cMap}(f, \text{groupBy}(T'(\mathcal{T}_c[\![c]\!]))) \\
&= \text{cMap}(f, \text{groupBy}(c))
\end{aligned}$$

If $e = \text{coGroup}(c_1, c_2)$, then using Equations (7h), (7f), (15b) and (15d)):

$$\begin{aligned}
&T'(\mathcal{T}_c[\![\text{cMap}(f, \text{coGroup}(c_1, c_2))]\!]) \\
&= T'(\text{sMap2}(f, \mathcal{T}_e[\![\text{coGroup}(c_1, c_2)]\!])) \\
&= T'(\text{sMap2}(f, \text{mix}(\text{coGroup}(\mathcal{T}_c[\![c_1]\!], \mathcal{T}_c[\![c_2]\!])))) \\
&= \text{cMap}(f, G_2(\text{mix}(\text{coGroup}(\mathcal{T}_c[\![c_1]\!], \mathcal{T}_c[\![c_2]\!])))) \\
&= \text{cMap}(f, \text{coGroup}(T'(\mathcal{T}_c[\![c_1]\!]), T'(\mathcal{T}_c[\![c_2]\!]))) \\
&= \text{cMap}(f, \text{coGroup}(c_1, c_2))
\end{aligned}$$

We now prove Equation (11) using induction on q. Based on Definition 4, we have three cases. When $q = (q_1, q_2)$, then for $x = \mathcal{T}_q[\![(q_1, q_2)]\!] = (\mathcal{T}_q[\![q_1]\!], \mathcal{T}_q[\![q_2]\!])$:

$$\begin{aligned}
&\mathcal{A}[\![(q_1, q_2)]\!]_x \\
&= (\mathcal{A}[\![q_1]\!]_{\pi_1(x)}, \mathcal{A}[\![q_2]\!]_{\pi_2(x)}) \\
&= (q_1, q_2) \qquad \text{(from induction hypothesis)}
\end{aligned}$$

For $q = \text{cMap}(f, S_i)$, we have:

$$\begin{aligned}
&\mathcal{A}[\![\text{cMap}(f, S_i)]\!]_x \\
&= \pi_2(x) \\
&= \pi_2(\mathcal{T}_q[\![\text{cMap}(f, S_i)]\!]) \\
&= \pi_2(\{\, ((), b) \mid a \in S_i,\, b \in f(a) \,\}) \\
&= \text{cMap}(f, S_i)
\end{aligned}$$

For $q = \text{cMap}(f, \text{groupBy}(c))$, we have:

$$\mathcal{A}[\![\text{cMap}(f, \text{groupBy}(c))]\!]_x = \pi_2(\text{reduce}(\Uparrow_{\otimes}, T(x)))$$

where

$$\begin{aligned}
x &= \mathcal{T}_q[\![\text{cMap}(f, \text{groupBy}(c))]\!] \\
&= \text{sMap1}(f, \text{groupBy}(\text{swap}(\mathcal{T}_c[\![c]\!])))
\end{aligned}$$

It is easy to prove by induction that for $X : \Uparrow_{\otimes}$ we have $\pi_2(\text{reduce}(\Uparrow_{\otimes}, T(\text{sMap1}(f, X)))) = \text{cMap}(f, G_1(X))$. Then,

$$\begin{aligned}
&\pi_2(\text{reduce}(\Uparrow_{\otimes}, T(\text{sMap1}(f, \text{groupBy}(\text{swap}(\mathcal{T}_c[\![c]\!])))))) \\
&= \text{cMap}(f, G_1(\text{groupBy}(\text{swap}(\mathcal{T}_c[\![c]\!])))) \\
&= \text{cMap}(f, \text{groupBy}(T(\mathcal{T}_c[\![c]\!]))) \\
&= \text{cMap}(f, \text{groupBy}(c))
\end{aligned}$$

The proof for $q = \text{cMap}(f, \text{coGroup}(c_1, c_2))$ is similar. For $q = \text{reduce}(\oplus, \text{cMap}(f, e))$, and for $e = \text{groupBy}(c)$ or $e = \text{coGroup}(c_1, c_2)$ (the proof is easier for $e = S_i$), we have:

$$\begin{aligned}
&\mathcal{A}[\![\text{reduce}(\oplus, \text{cMap}(f, e))]\!]_x \\
&= \text{reduce}(\oplus, \pi_2(x)) \\
&= \text{reduce}(\oplus, \pi_2(\mathcal{T}_q[\![\text{reduce}(\oplus, \text{cMap}(f, e))]\!])) \\
&= \text{reduce}(\oplus, \pi_2(\text{reduce}(\Uparrow_{\oplus}, \text{sMap1}(f, \mathcal{T}_e[\![e]\!])))) \\
&= \text{reduce}(\oplus, \pi_2(\text{reduce}(\Uparrow_{\oplus}, \mathcal{T}_q[\![\text{cMap}(f, e)]\!]))) \\
&= \text{reduce}(\oplus, \text{cMap}(f, e))
\end{aligned}$$

from the previous proofs for $q = \text{cMap}(f, \text{groupBy}(c))$ and $q = \text{cMap}(f, \text{coGroup}(c_1, c_2))$. □

Theorem 3. *For all $X, Y : (X \otimes Y) \mathbin{\widetilde{\otimes}} Y = X$.*

Proof. We will prove this for $\Downarrow_{\widetilde{\oplus}}$ only:

$$\begin{aligned}
&(X \Uparrow_{\oplus} Y) \Downarrow_{\widetilde{\oplus}} Y \\
&= \{\, (k, z \mathbin{\widetilde{\oplus}} y) \mid (k, z) \in (X \Uparrow_{\oplus} Y),\, (k, y) \in Y,\, z \neq y \,\} \\
&\quad \uplus\ \{\, (k, z) \mid (k, z) \in (X \Uparrow_{\oplus} Y),\, k \notin \pi_1(Y) \,\} \\
&= \{\, (k, (x \oplus y) \mathbin{\widetilde{\oplus}} y) \mid (k, x) \in X,\, (k, y) \in Y, \\
&\qquad\qquad (k, y) \in Y,\, (x \oplus y) \neq y \,\} \\
&\quad \uplus\ \{\, (k, z) \mid (k, z) \in (\{\, (k, x \oplus y) \mid (k, x) \in X,\, (k, y) \in Y \,\} \\
&\qquad\qquad \uplus\ \{\, (k, x) \mid (k, x) \in X,\, k \notin \pi_1(Y) \,\} \\
&\qquad\qquad \uplus\ \{\, (k, y) \mid (k, y) \in Y,\, k \notin \pi_1(X) \,\}), \\
&\qquad k \notin \pi_1(Y) \,\} \\
&= \{\, (k, (x \oplus y) \mathbin{\widetilde{\oplus}} y) \mid (k, x) \in X,\, (k, y) \in Y, \\
&\qquad\qquad (k, y) \in Y,\, (x \oplus y) \neq y \,\} \\
&\quad \uplus\ \{\, (k, x) \mid (k, x) \in X,\, k \notin \pi_1(Y),\, k \notin \pi_1(Y) \,\} \\
&= \{\, (k, x) \mid (k, x) \in X,\, (k, y) \in Y \,\} \\
&\quad \uplus\ \{\, (k, x) \mid (k, x) \in X,\, k \notin \pi_1(Y) \,\} \\
&= X
\end{aligned}$$

□